



(12) **United States Patent**
Busaba et al.

(10) **Patent No.:** US 11,182,198 B2
(45) **Date of Patent:** *Nov. 23, 2021

(54) **INDICATOR-BASED PRIORITIZATION OF TRANSACTIONS**

(52) **U.S. Cl.**
CPC **G06F 9/466** (2013.01); **G06F 9/467** (2013.01); **G06F 9/5038** (2013.01);
(Continued)

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

(58) **Field of Classification Search**
None
See application file for complete search history.

(72) Inventors: **Fadi Y. Busaba**, Poughkeepsie, NY (US); **Harold W. Cain, III**, Raleigh, NC (US); **Michael K. Gschwind**, Chappaqua, NY (US); **Valentina Salapura**, Chappaqua, NY (US); **Eric M. Schwarz**, Gardiner, NY (US); **Timothy J. Siegel**, Staatsburg, NY (US)

(56) **References Cited**
U.S. PATENT DOCUMENTS

7,571,284 B1 8/2009 Olson
7,711,909 B1 5/2010 Lev et al.
(Continued)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 89 days.

Bardhan et al., "Predicting the Effect of Memory Contention in Multi-Core Computers Using Analytic Performance Models, 2015, IEEE Transaction on Computers, vol. 64, No. 8, pp. 2279-2292.

This patent is subject to a terminal disclaimer.

(Continued)

Primary Examiner — Abu Zar Ghaffari
(74) *Attorney, Agent, or Firm* — Heslin Rothenberg Farley & Mesiti P.C.; Steven Chiu, Esq.

(21) Appl. No.: **16/415,502**

(22) Filed: **May 17, 2019**

(65) **Prior Publication Data**

US 2019/0272192 A1 Sep. 5, 2019

Related U.S. Application Data

(63) Continuation of application No. 15/277,770, filed on Sep. 27, 2016, now Pat. No. 10,353,734, which is a
(Continued)

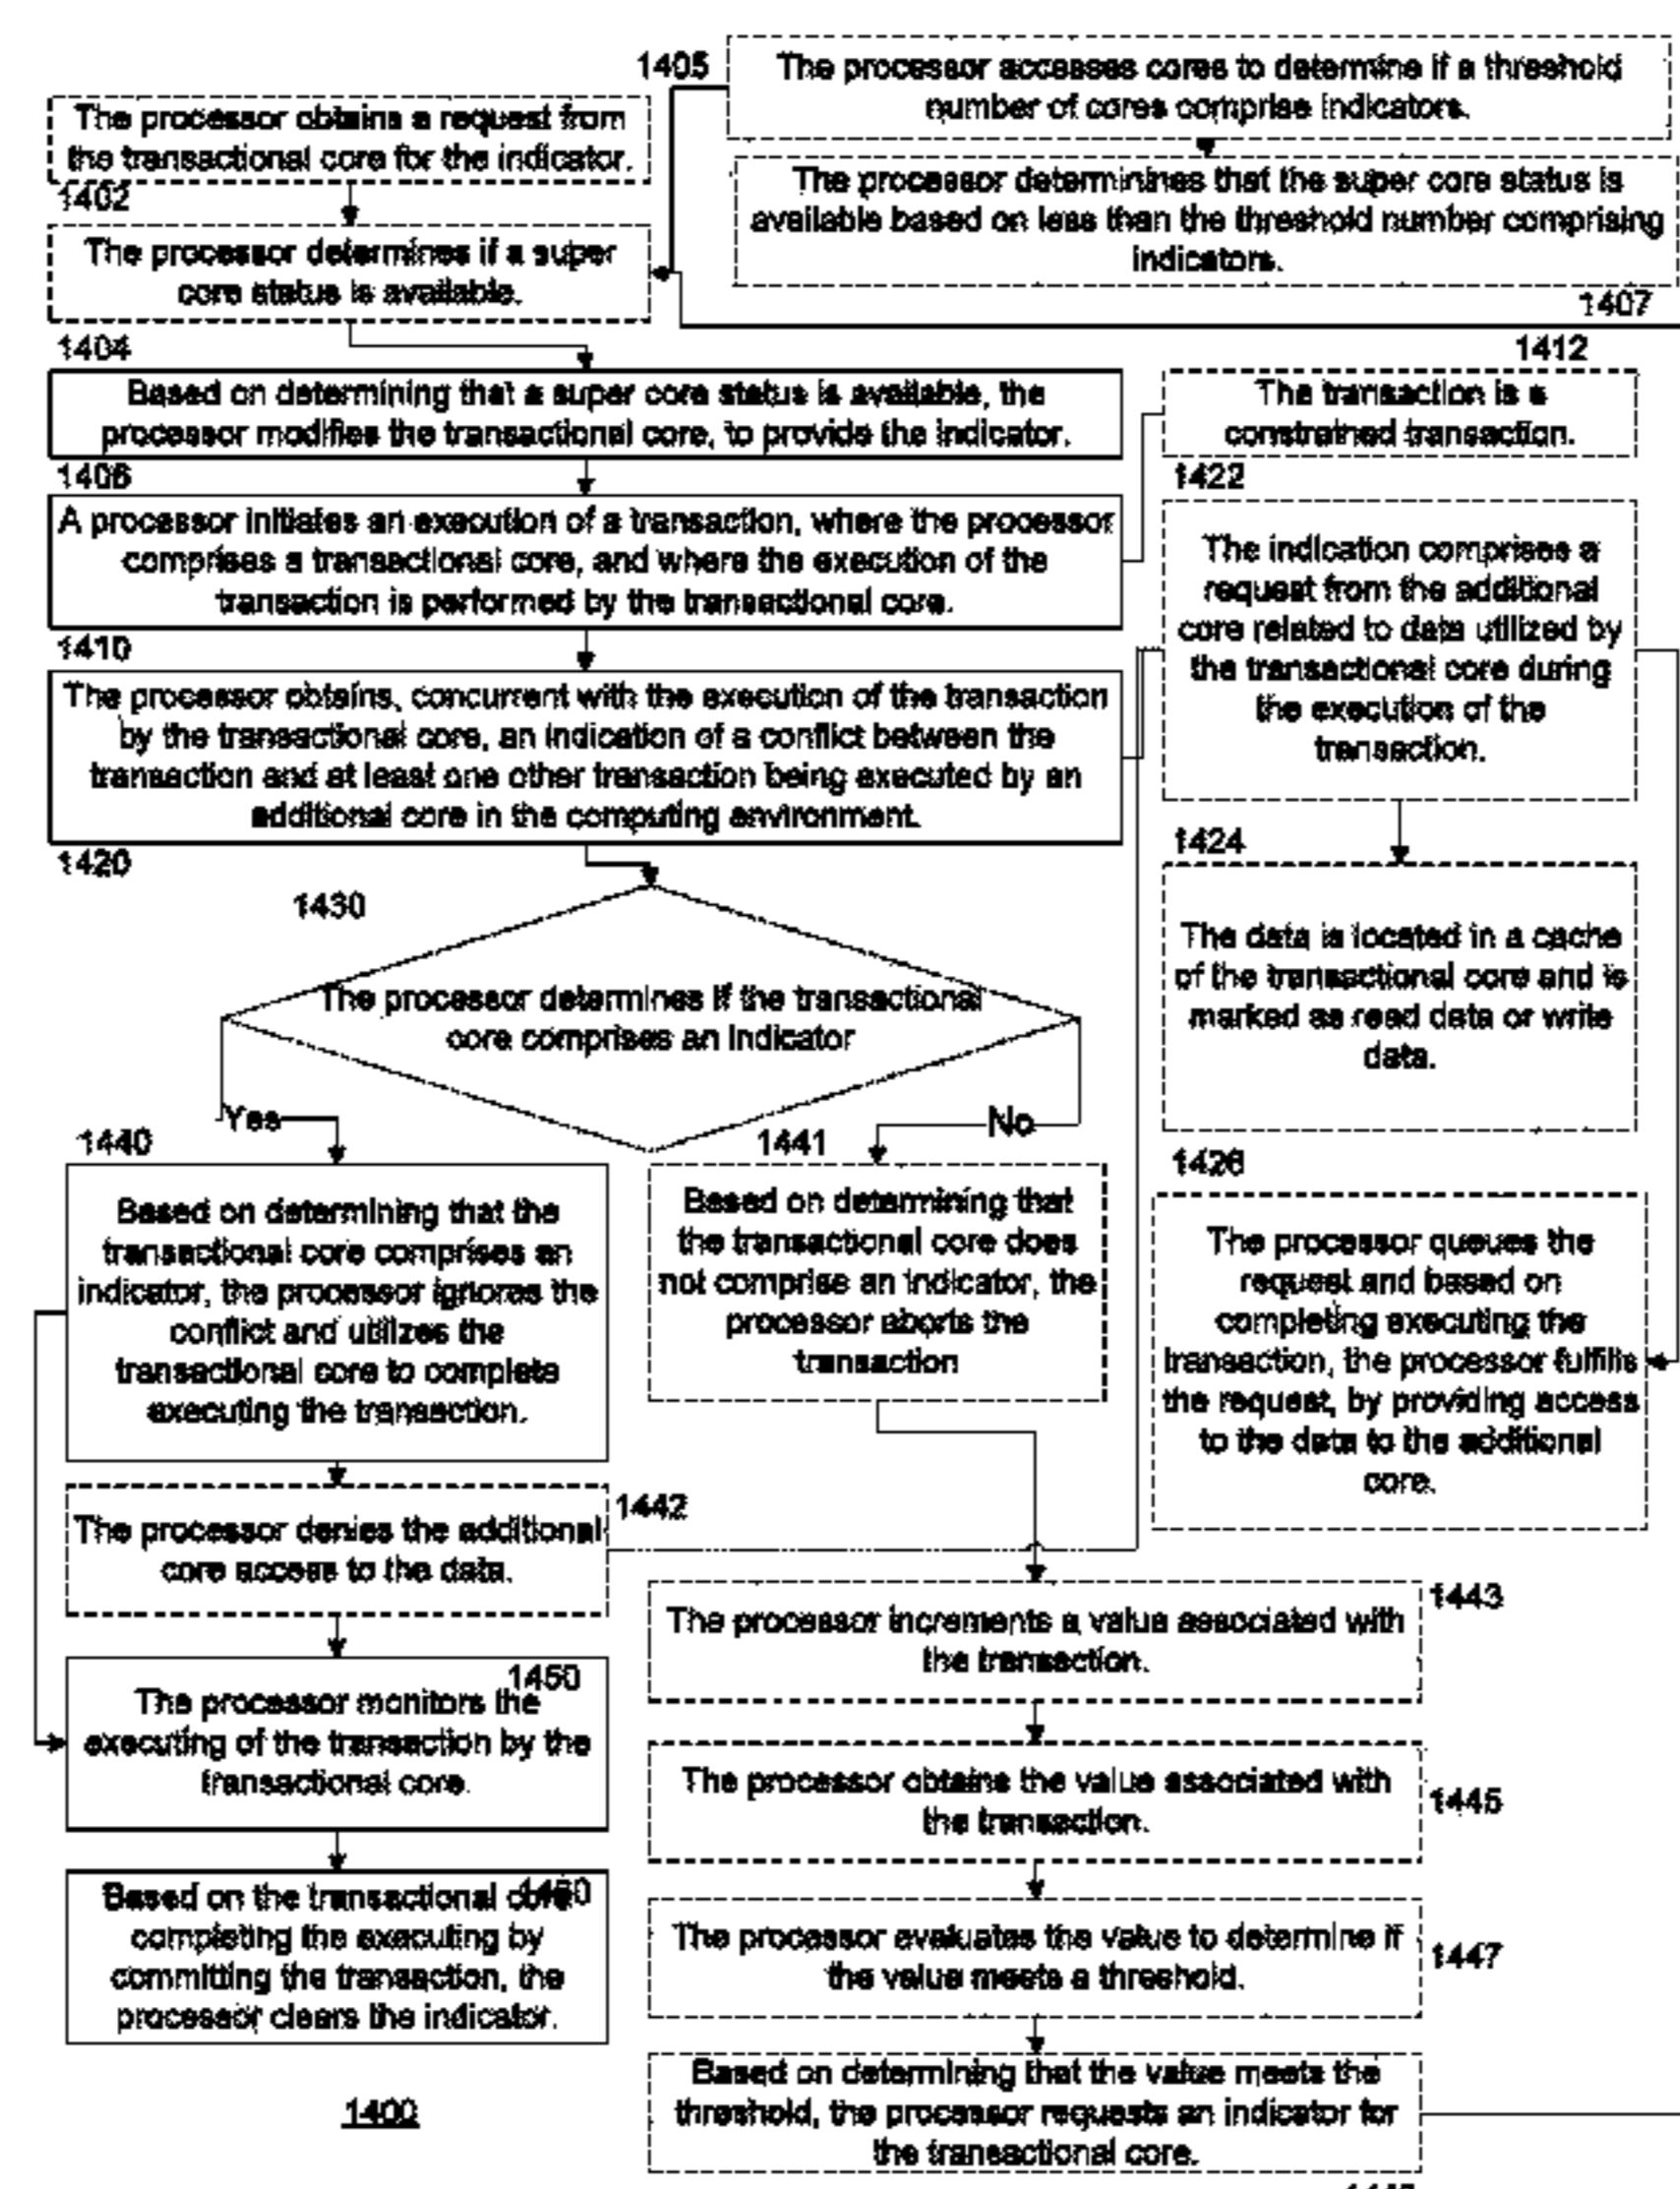
(51) **Int. Cl.**
G06F 9/46 (2006.01)
G06F 12/0831 (2016.01)

(Continued)

(57) **ABSTRACT**

A method, system, and computer program product are provided for prioritizing transactions. A processor in a computing environment initiates the execution of a transaction. The processor includes a transactional core, and the execution of the transaction is performed by the transactional core. The processor obtains concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional core in the computing environment. The processor determines if the transactional core includes an indicator and based on determining that the transactional core includes an indicator, the processor ignores the conflict and utilizing the transactional core to complete executing the transaction.

16 Claims, 24 Drawing Sheets



Related U.S. Application Data

continuation of application No. 15/010,681, filed on Jan. 29, 2016, now Pat. No. 9,772,874.

(51) **Int. Cl.**

G06F 12/0817 (2016.01)
G06F 9/50 (2006.01)
G06F 12/02 (2006.01)
G06F 9/30 (2018.01)
G06F 9/52 (2006.01)
G06F 12/14 (2006.01)

(52) **U.S. Cl.**

CPC **G06F 12/0828** (2013.01); **G06F 12/0831** (2013.01); *G06F 9/3004* (2013.01); *G06F 9/30087* (2013.01); *G06F 9/30189* (2013.01); *G06F 9/526* (2013.01); *G06F 12/0284* (2013.01); *G06F 12/1416* (2013.01); *G06F 2212/1016* (2013.01); *G06F 2212/1032* (2013.01); *G06F 2212/1052* (2013.01); *G06F 2212/152* (2013.01); *G06F 2212/621* (2013.01)

(56)

References Cited

U.S. PATENT DOCUMENTS

7,730,261 B1 6/2010 Yung et al.
 8,037,476 B1 10/2011 Shavit et al.
 8,078,847 B2 12/2011 Gonion et al.
 8,108,625 B1 1/2012 Coon et al.
 8,356,166 B2* 1/2013 Adl-Tabatabai G06F 9/30181
 712/229
 8,358,651 B1 1/2013 Kadosh
 8,365,016 B2 1/2013 Gray et al.
 8,886,918 B2 11/2014 Capps, Jr. et al.
 9,002,872 B2 4/2015 Jose
 9,244,746 B2* 1/2016 Hughes G06F 9/467
 9,268,542 B1 2/2016 Mars
 9,390,034 B1 7/2016 Fair
 9,465,387 B2 10/2016 Ross
 9,772,874 B2 9/2017 Busaba et al.
 2003/0145146 A1 7/2003 Posner et al.
 2004/0154018 A1 8/2004 Doerling
 2004/0170046 A1 9/2004 Belnet et al.
 2005/0154861 A1 7/2005 Arimilli
 2007/0156994 A1 7/2007 Akkary et al.
 2008/0163255 A1 7/2008 Munoz et al.
 2008/0320476 A1 12/2008 Wingard et al.
 2009/0070774 A1 3/2009 Raikin et al.
 2009/0077329 A1* 3/2009 Wood G06F 12/0815
 711/156

2009/0100255 A1 4/2009 Harris
 2009/0138890 A1 5/2009 Blake et al.
 2009/0187906 A1 7/2009 Caprioli et al.
 2009/0198694 A1* 8/2009 Thomas G06F 9/466
 2010/0107174 A1 4/2010 Suzuki et al.
 2010/0174840 A1 7/2010 Blainey et al.
 2011/0078385 A1 3/2011 Lev et al.
 2011/0209151 A1* 8/2011 Chung G06F 9/3834
 718/101
 2011/0209155 A1 8/2011 Giampapa et al.
 2011/0246993 A1 10/2011 Moir et al.
 2012/0210162 A1 8/2012 Gara et al.
 2013/0117754 A1 5/2013 Kurihara et al.
 2013/0198478 A1 8/2013 Bitner
 2013/0339676 A1 12/2013 Greiner et al.
 2013/0339703 A1 12/2013 Alexander et al.
 2014/0040551 A1 2/2014 Blainey et al.
 2014/0040567 A1* 2/2014 Pohlack G06F 9/524
 711/152
 2015/0100741 A1 4/2015 Calciu et al.
 2015/0100753 A1 4/2015 Shen et al.
 2015/0212851 A1 7/2015 Busaba et al.
 2015/0277967 A1 10/2015 Calciu et al.
 2015/0378927 A1 12/2015 Bradbury et al.
 2017/0046198 A1 2/2017 Solihin
 2017/0220470 A1 8/2017 Busaba et al.

OTHER PUBLICATIONS

Blake et al., "Proactive Transaction Scheduling for Contention Management", MICRO Dec. 12-16, 2009, pp. 156-167.
 Spear et al., "A Comprehensive Strategy for Contention Management in Software Transactional Memory", PPOPP' 09, Feb. 14-18, 2009, pp. 1-10.
 Rakesh Agrawal' "Concurrency Control and Recovery in Multiprocessor Database Machines: Design and Performance Evaluation", IP.com Electronic Publication, Dec. 8, 2007, IP.com No. IPCOM00161111D, 214 pages.
 Njuguna Njoroge et al., "ATLAS: A chip Multiprocessor With Transaction Memory Support", Computer Systems Laboratory, Stanford University, 6 pages.
 Mell et al., "The NIST Definition Of Cloud Computing", National Institute of Standards and Technology, U.S. Department of Commerce, NIST Special Publication 800-145, Sep. 2011, 7 pages.
 International Search Report and the Written Opinion, International Application No. PCT/IB2017/050147, dated May 11, 2017, 11 pages.
 List of IBM Patents or Patent Applications Treated as Related, Jun. 14, 2019, 2 pages.

* cited by examiner

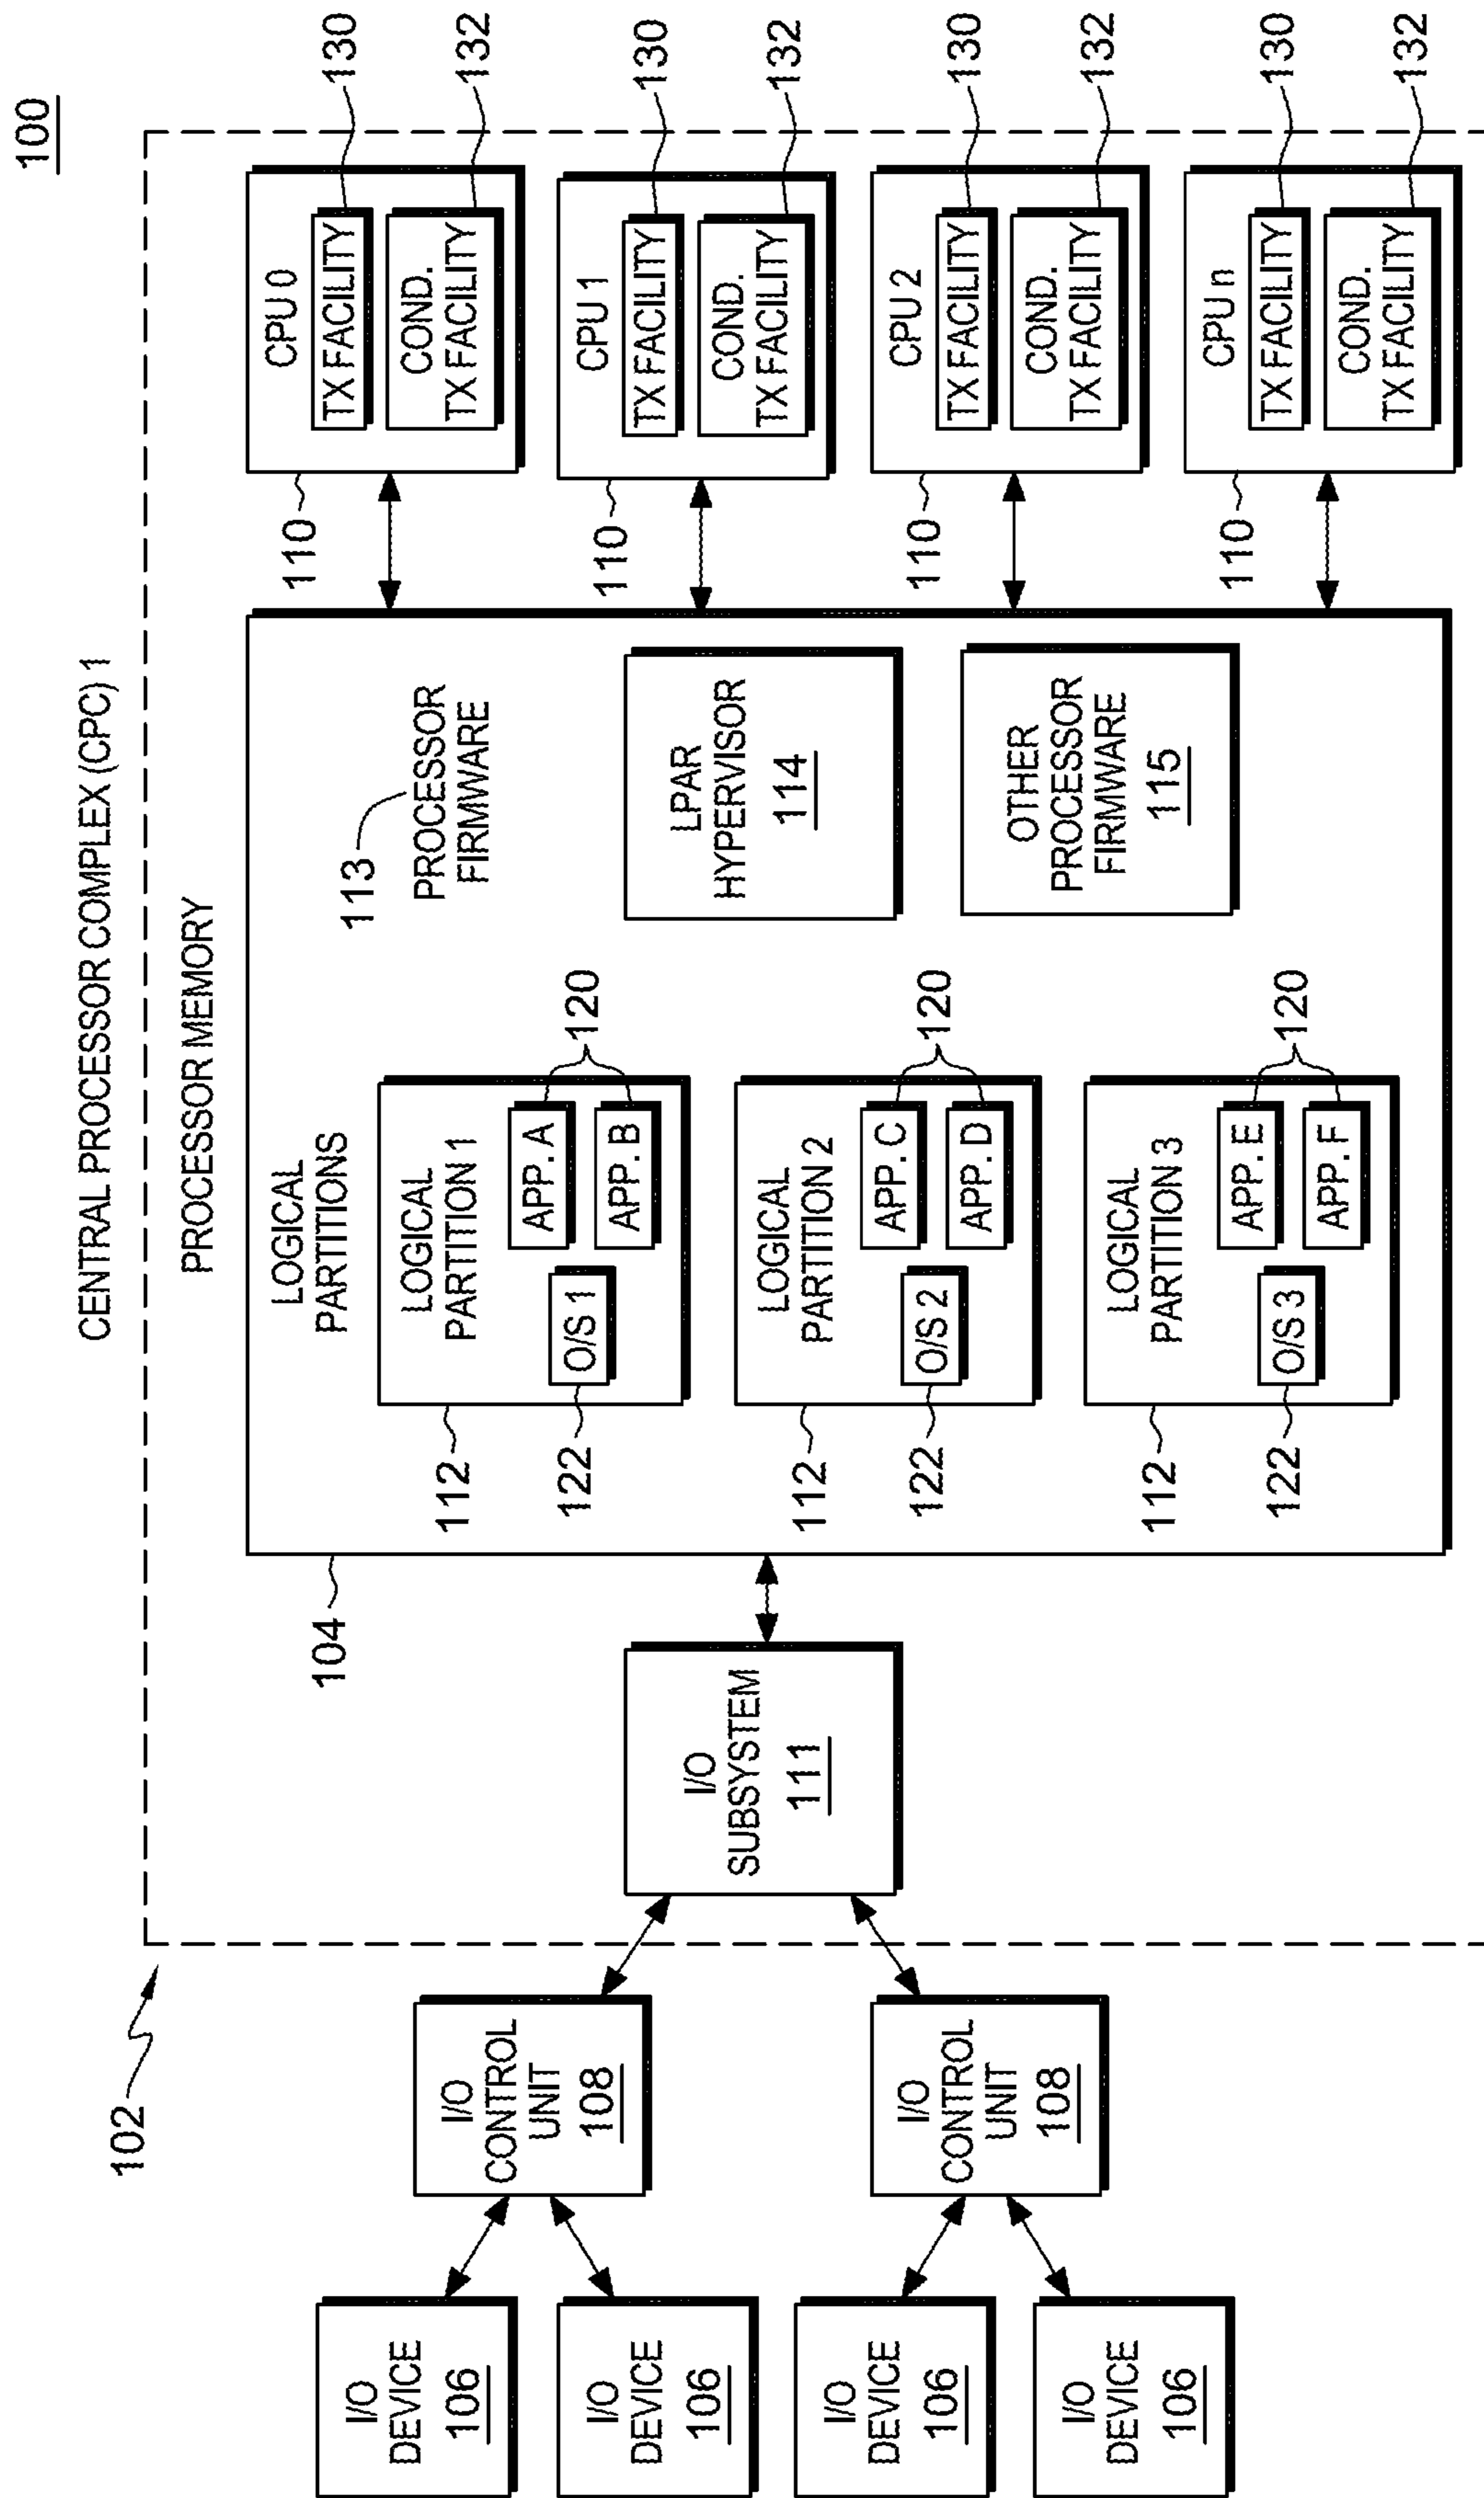


FIG. 1

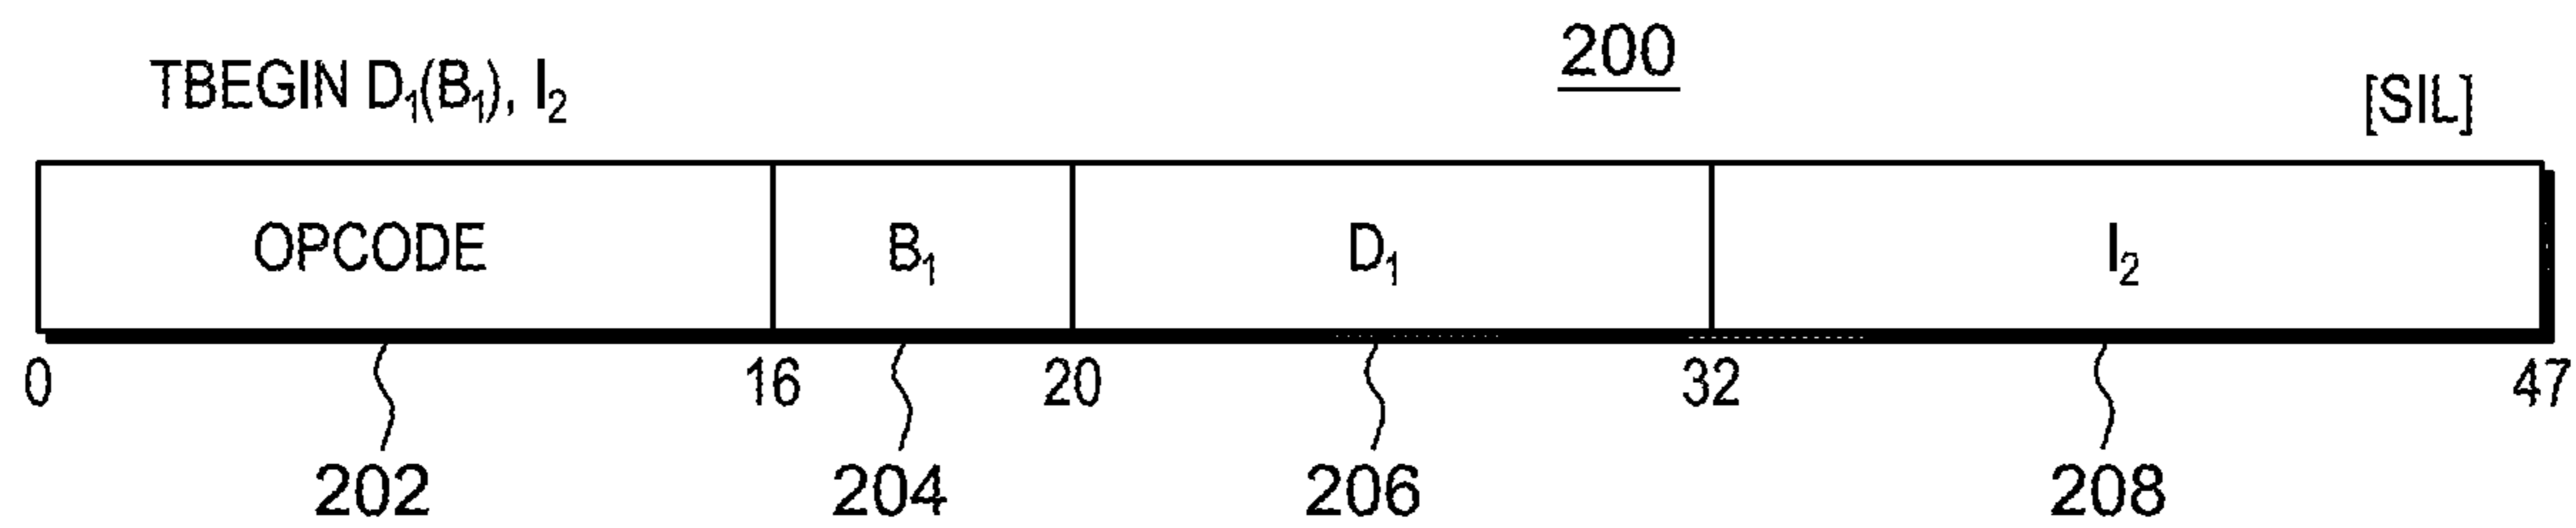


FIG. 2A

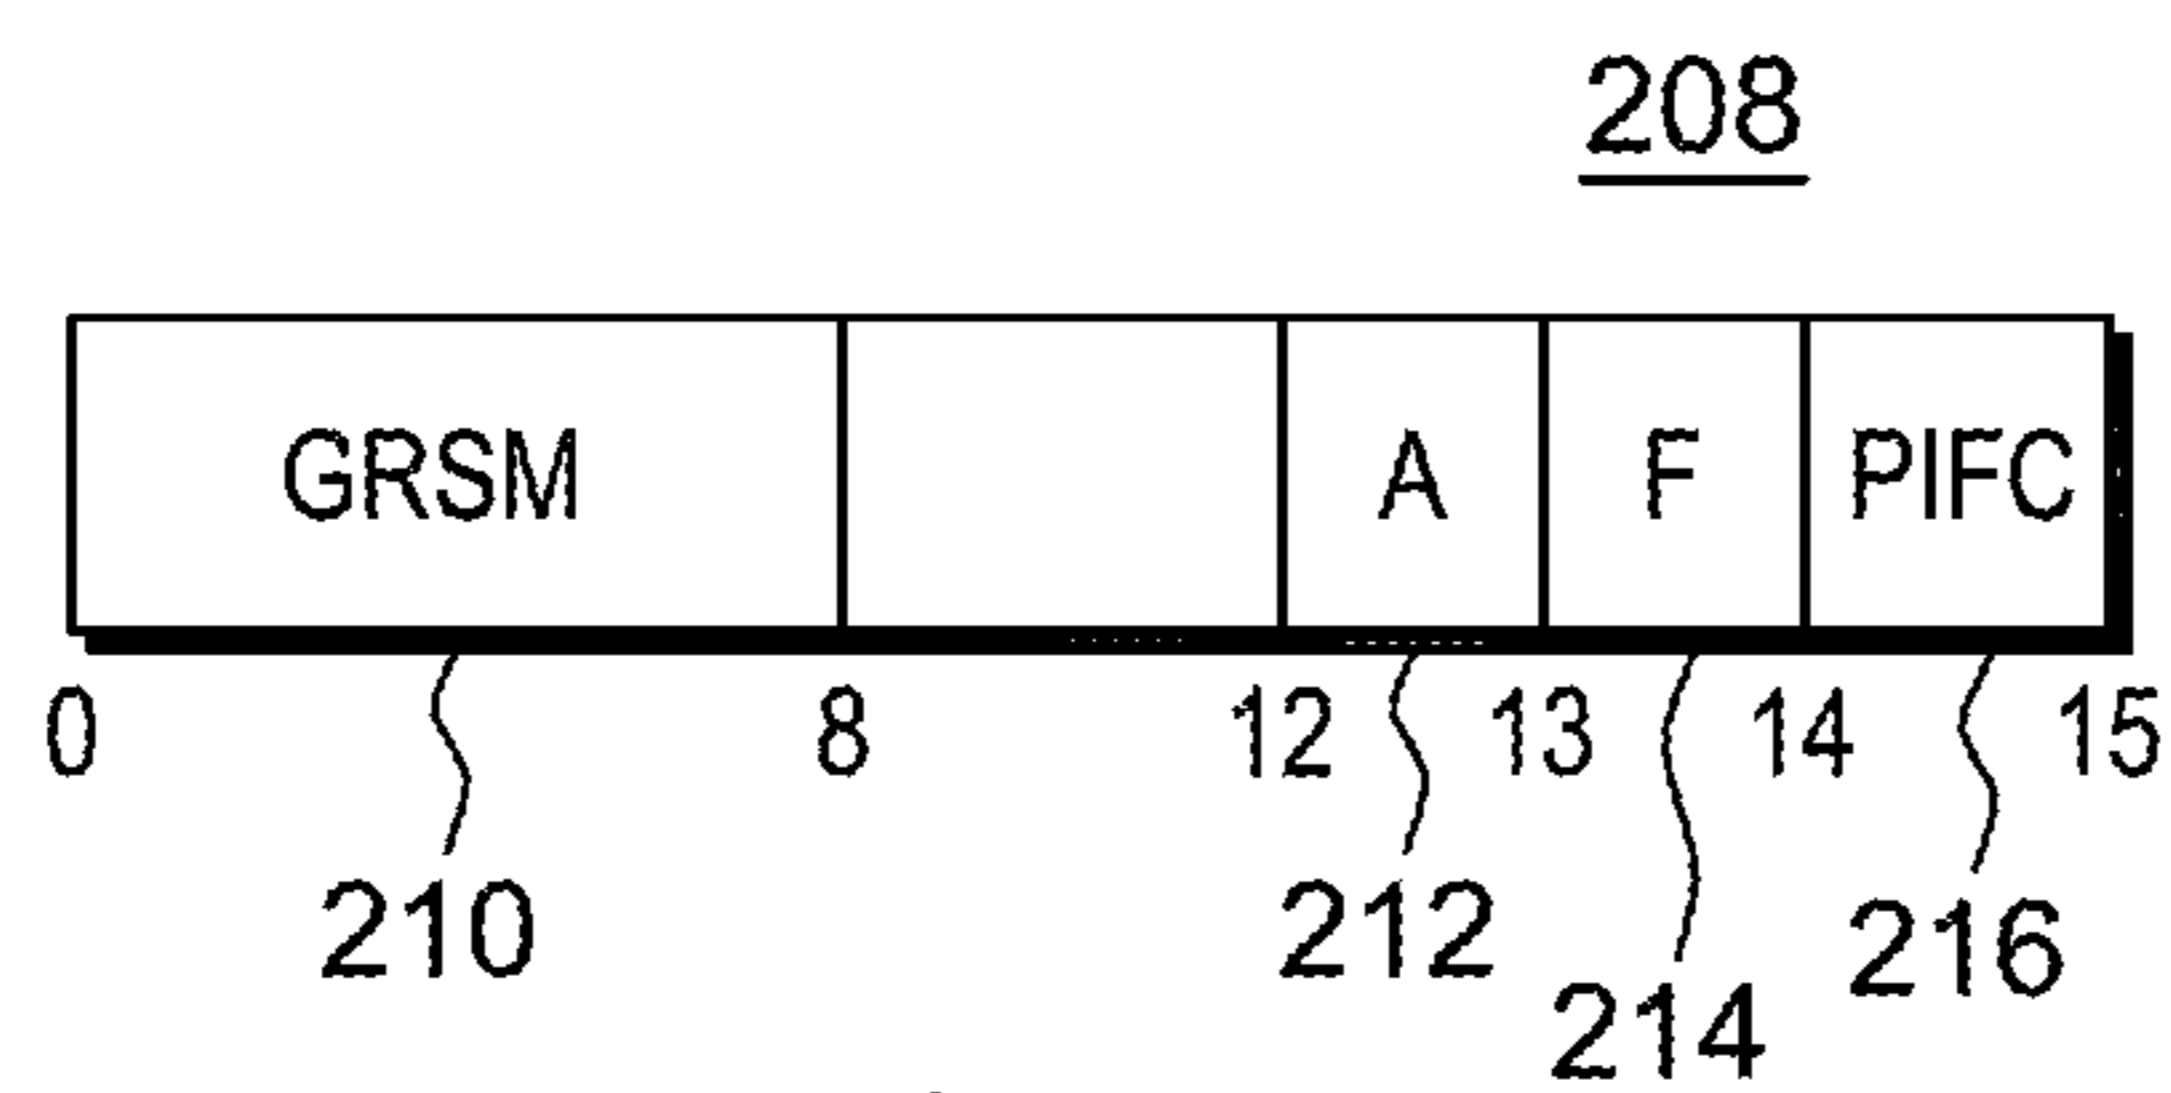


FIG. 2B

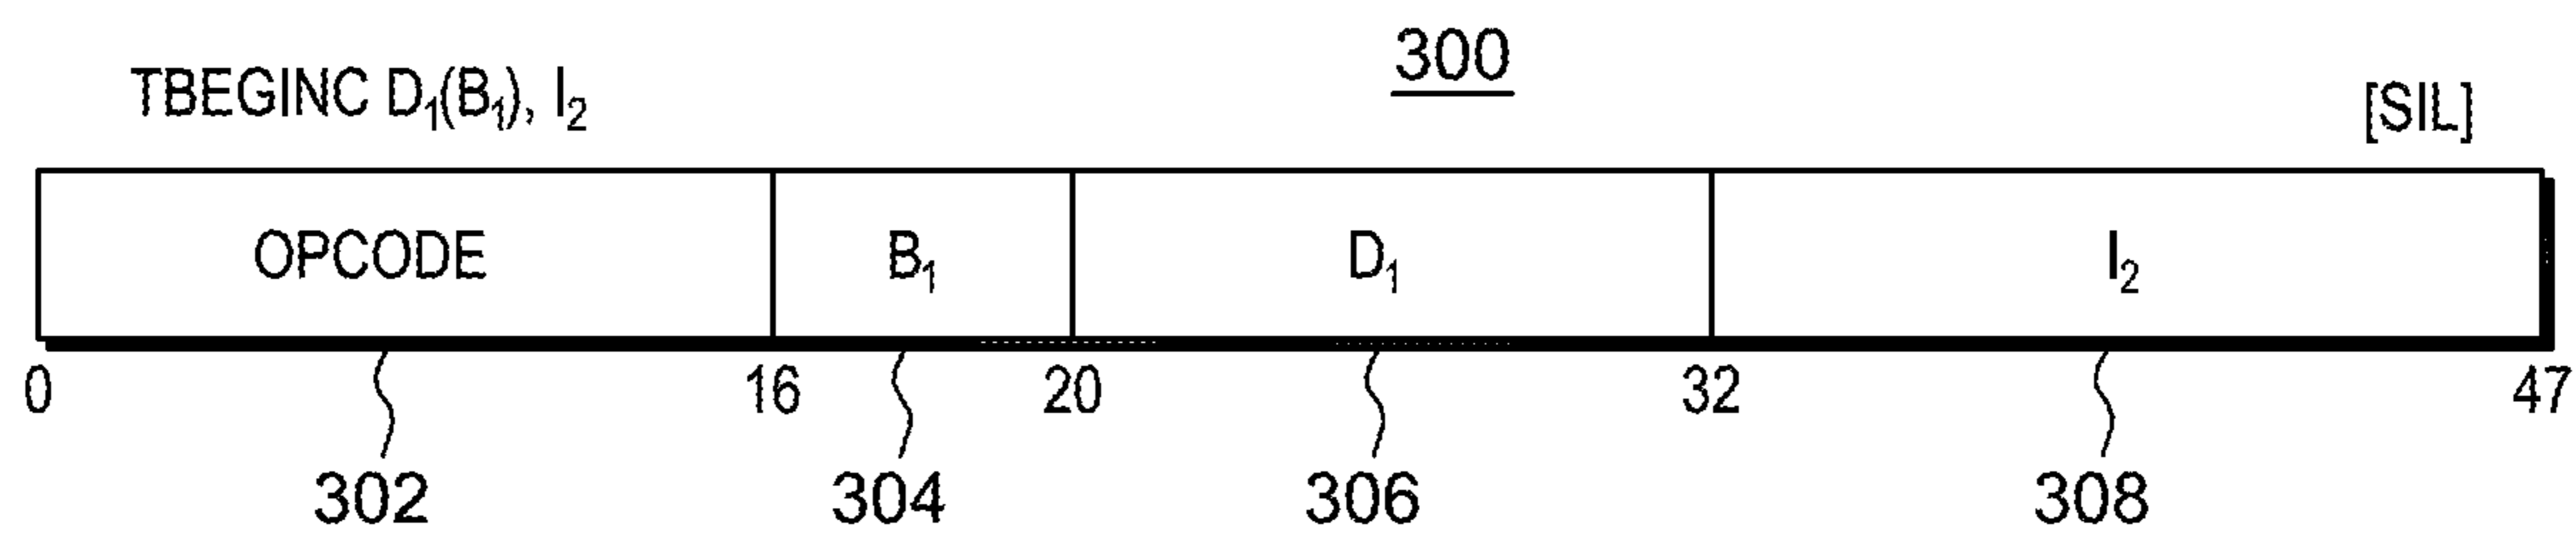


FIG. 3A

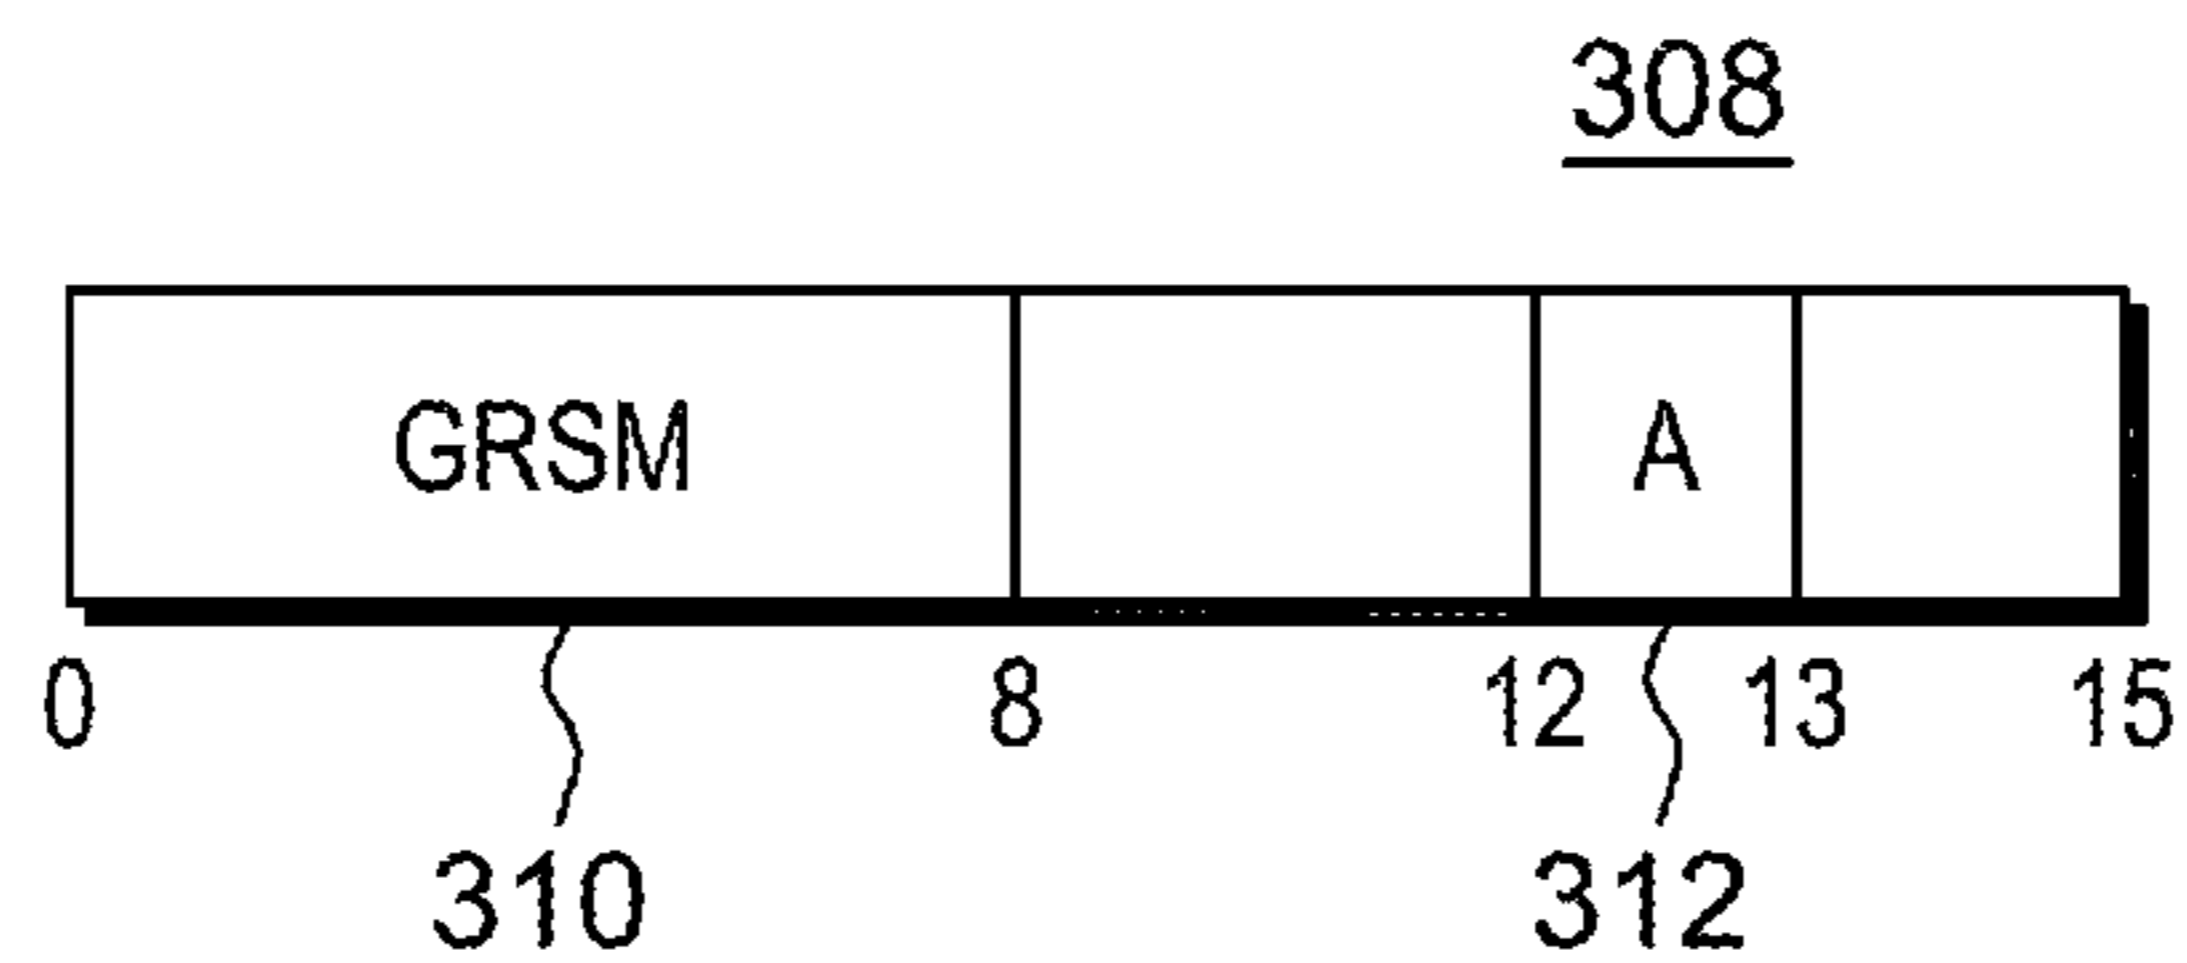


FIG. 3B

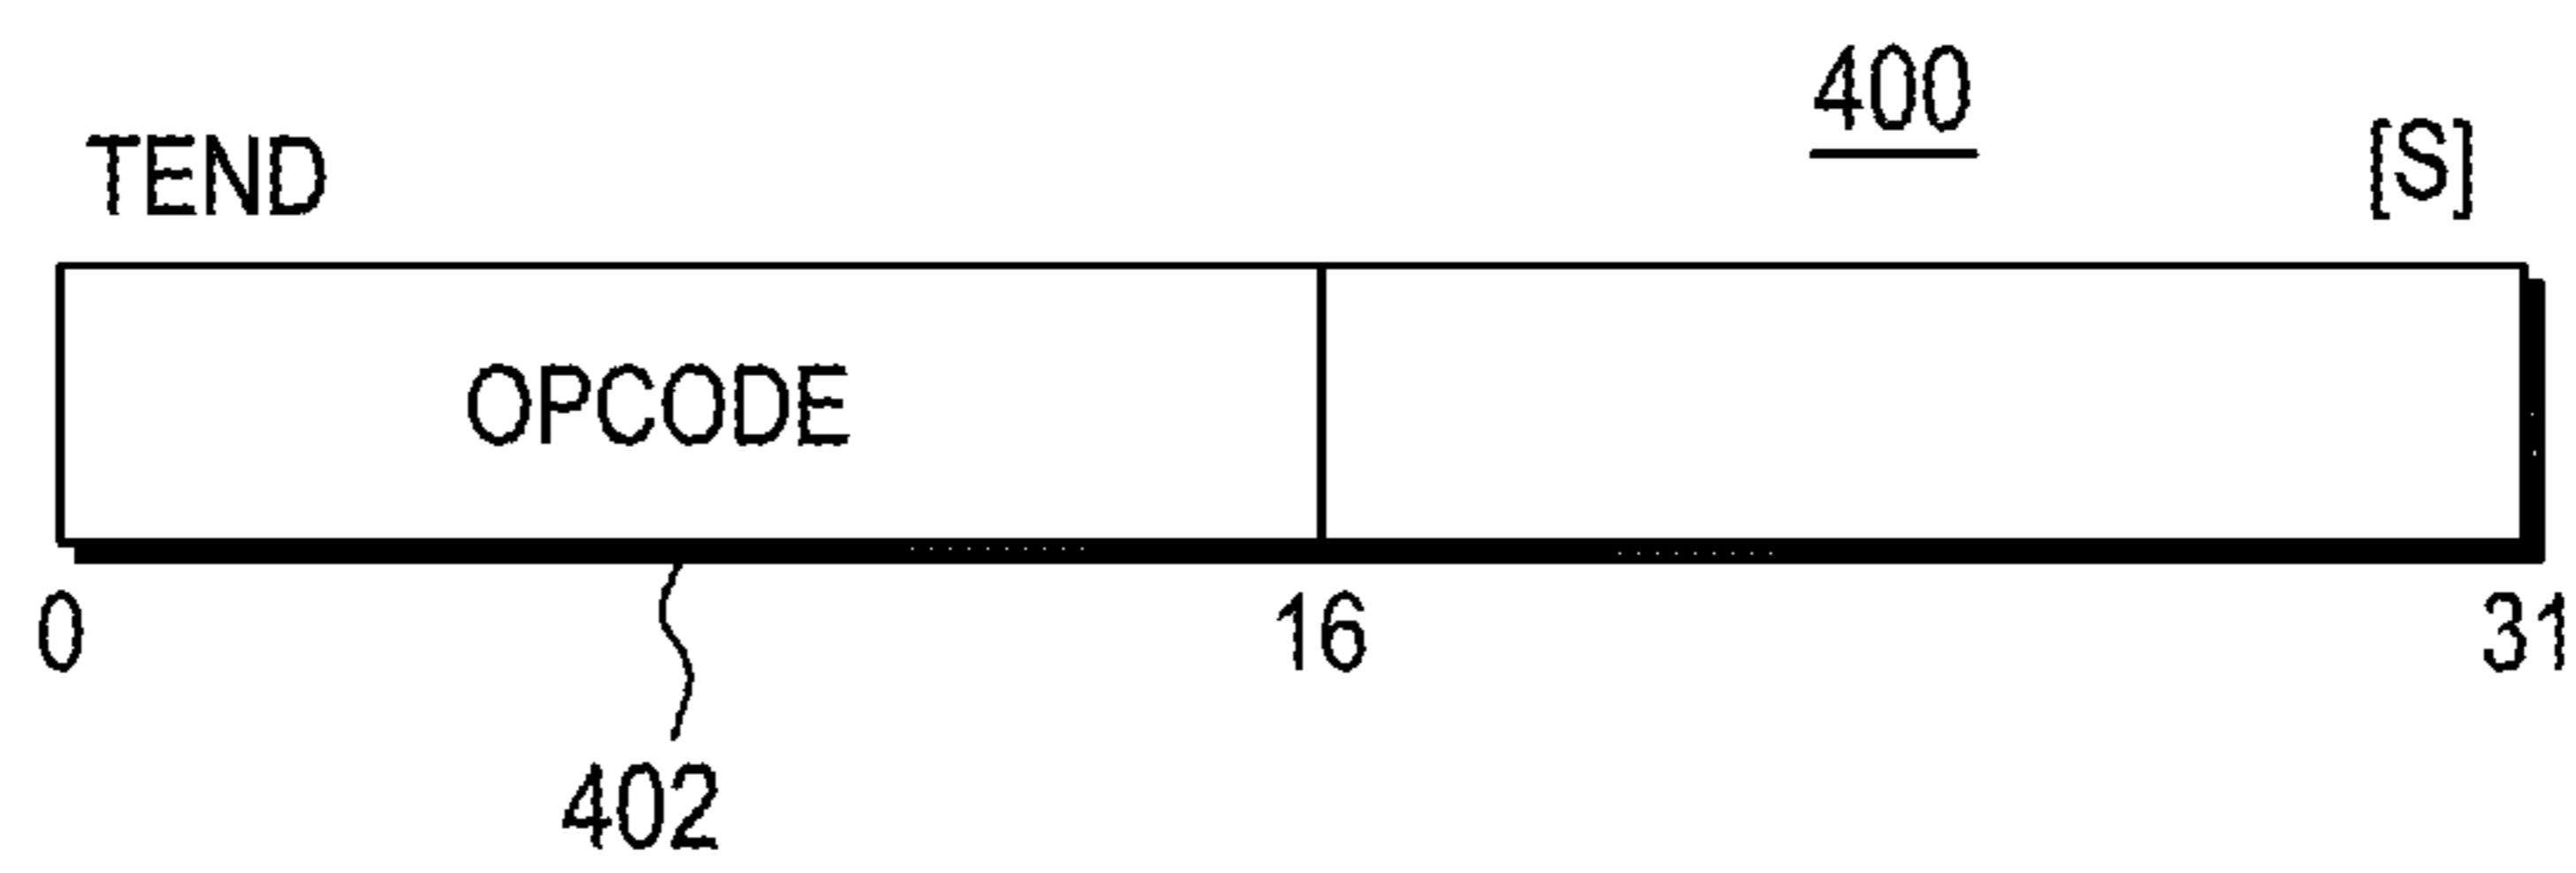


FIG. 4

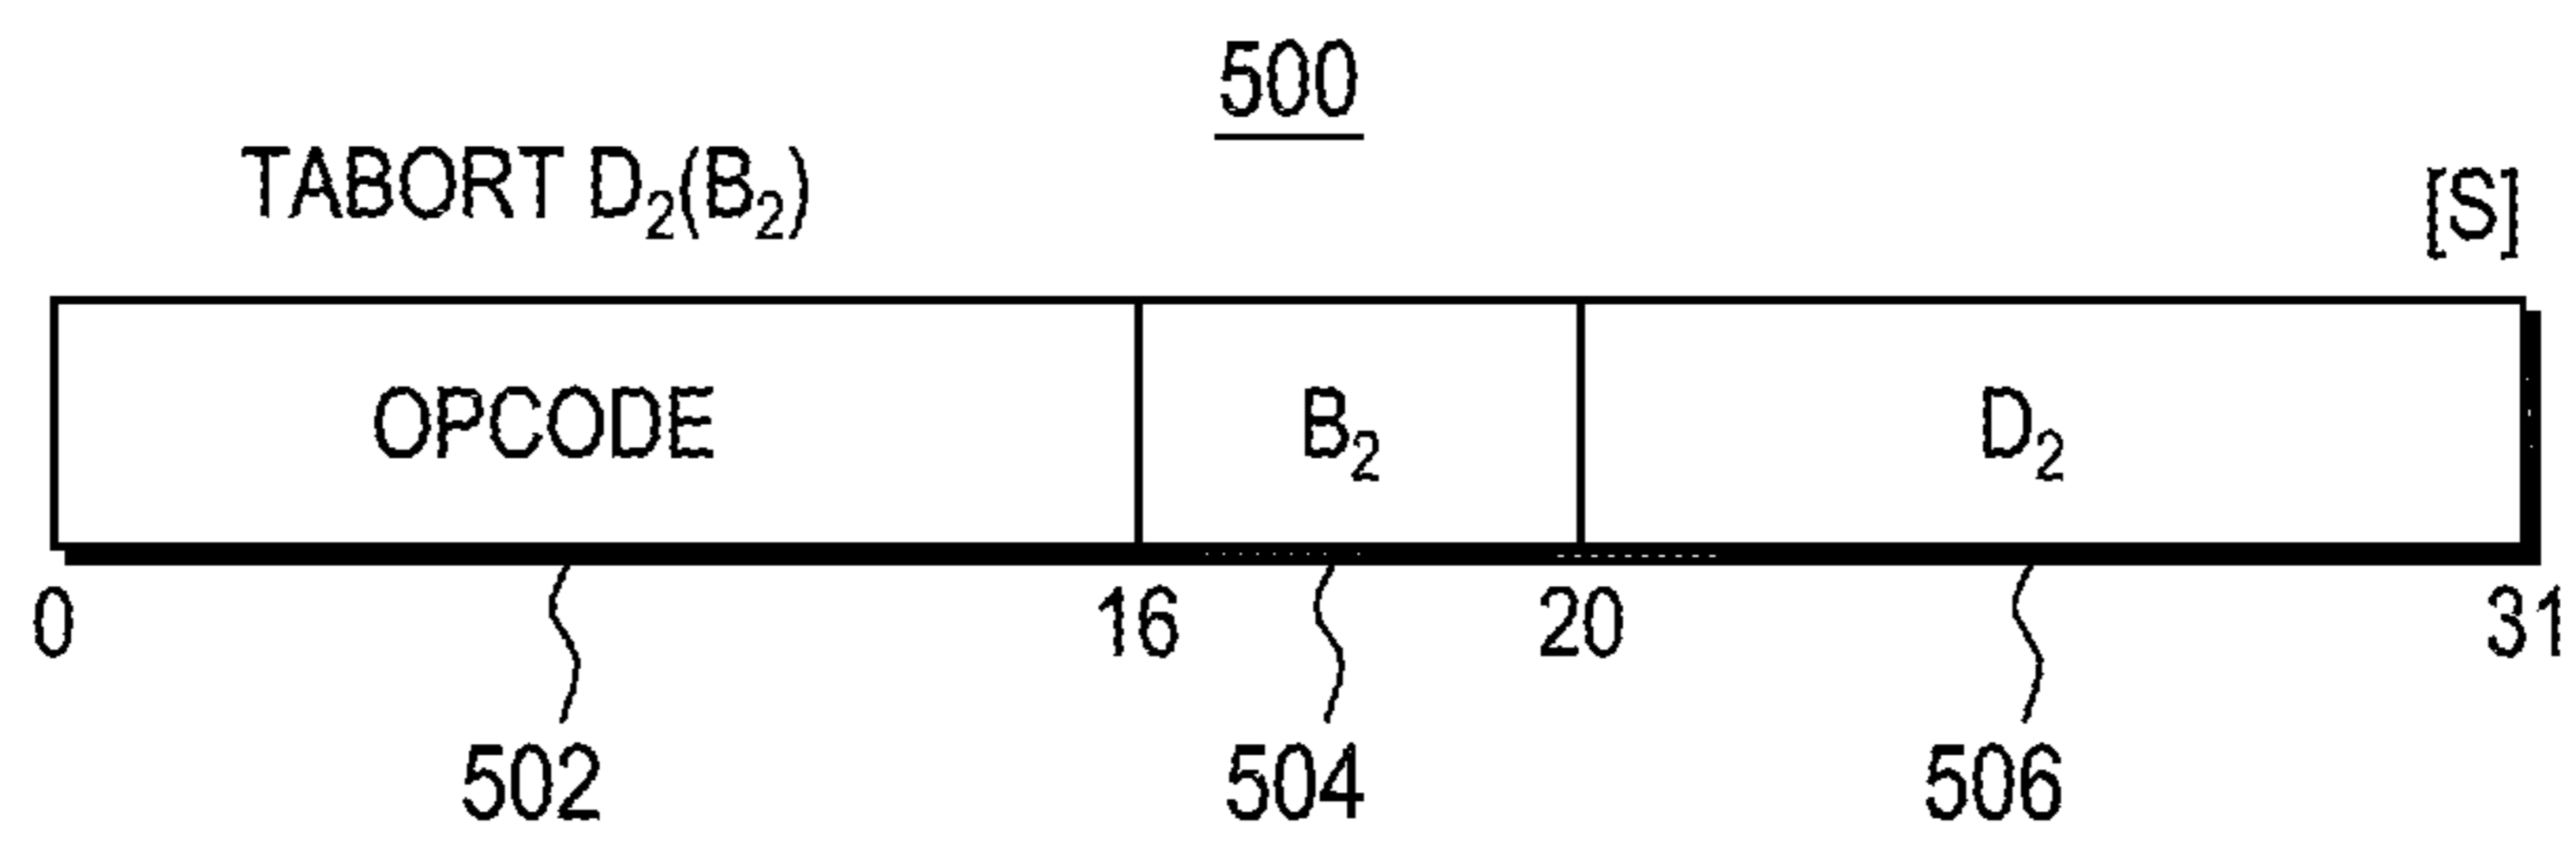


FIG. 5

NESTED TRANSACTIONS

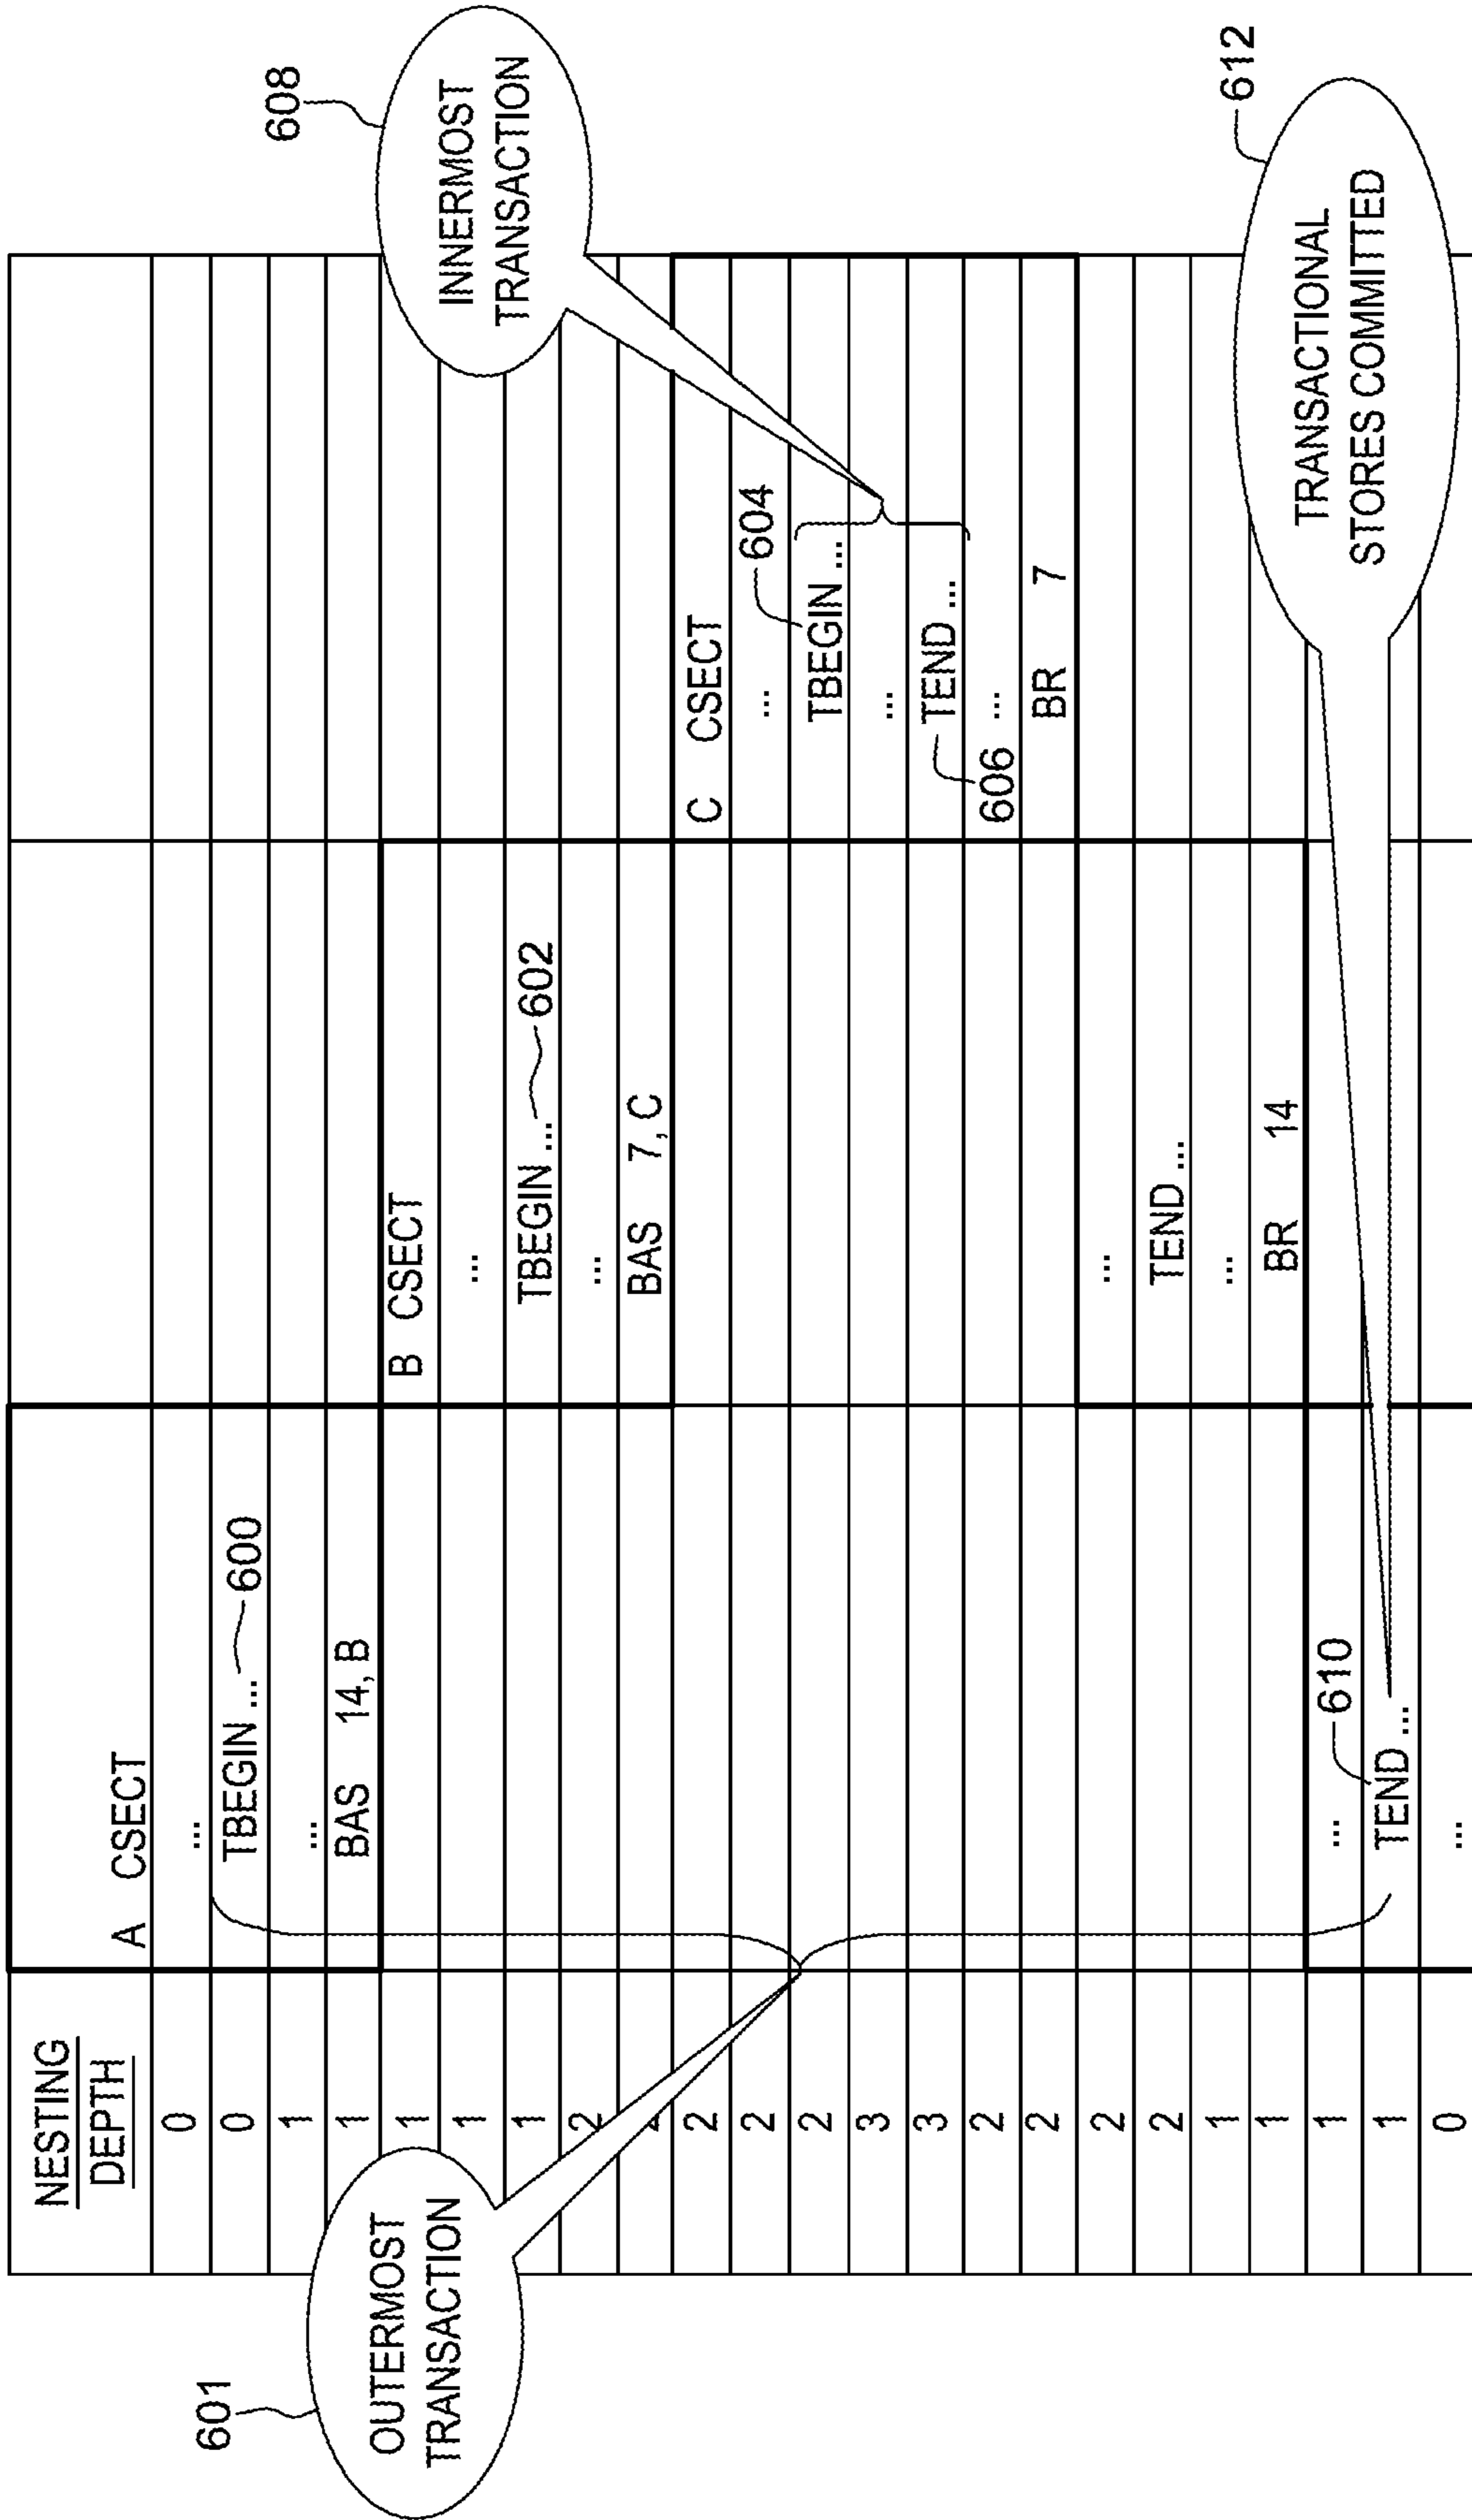


FIG. 6

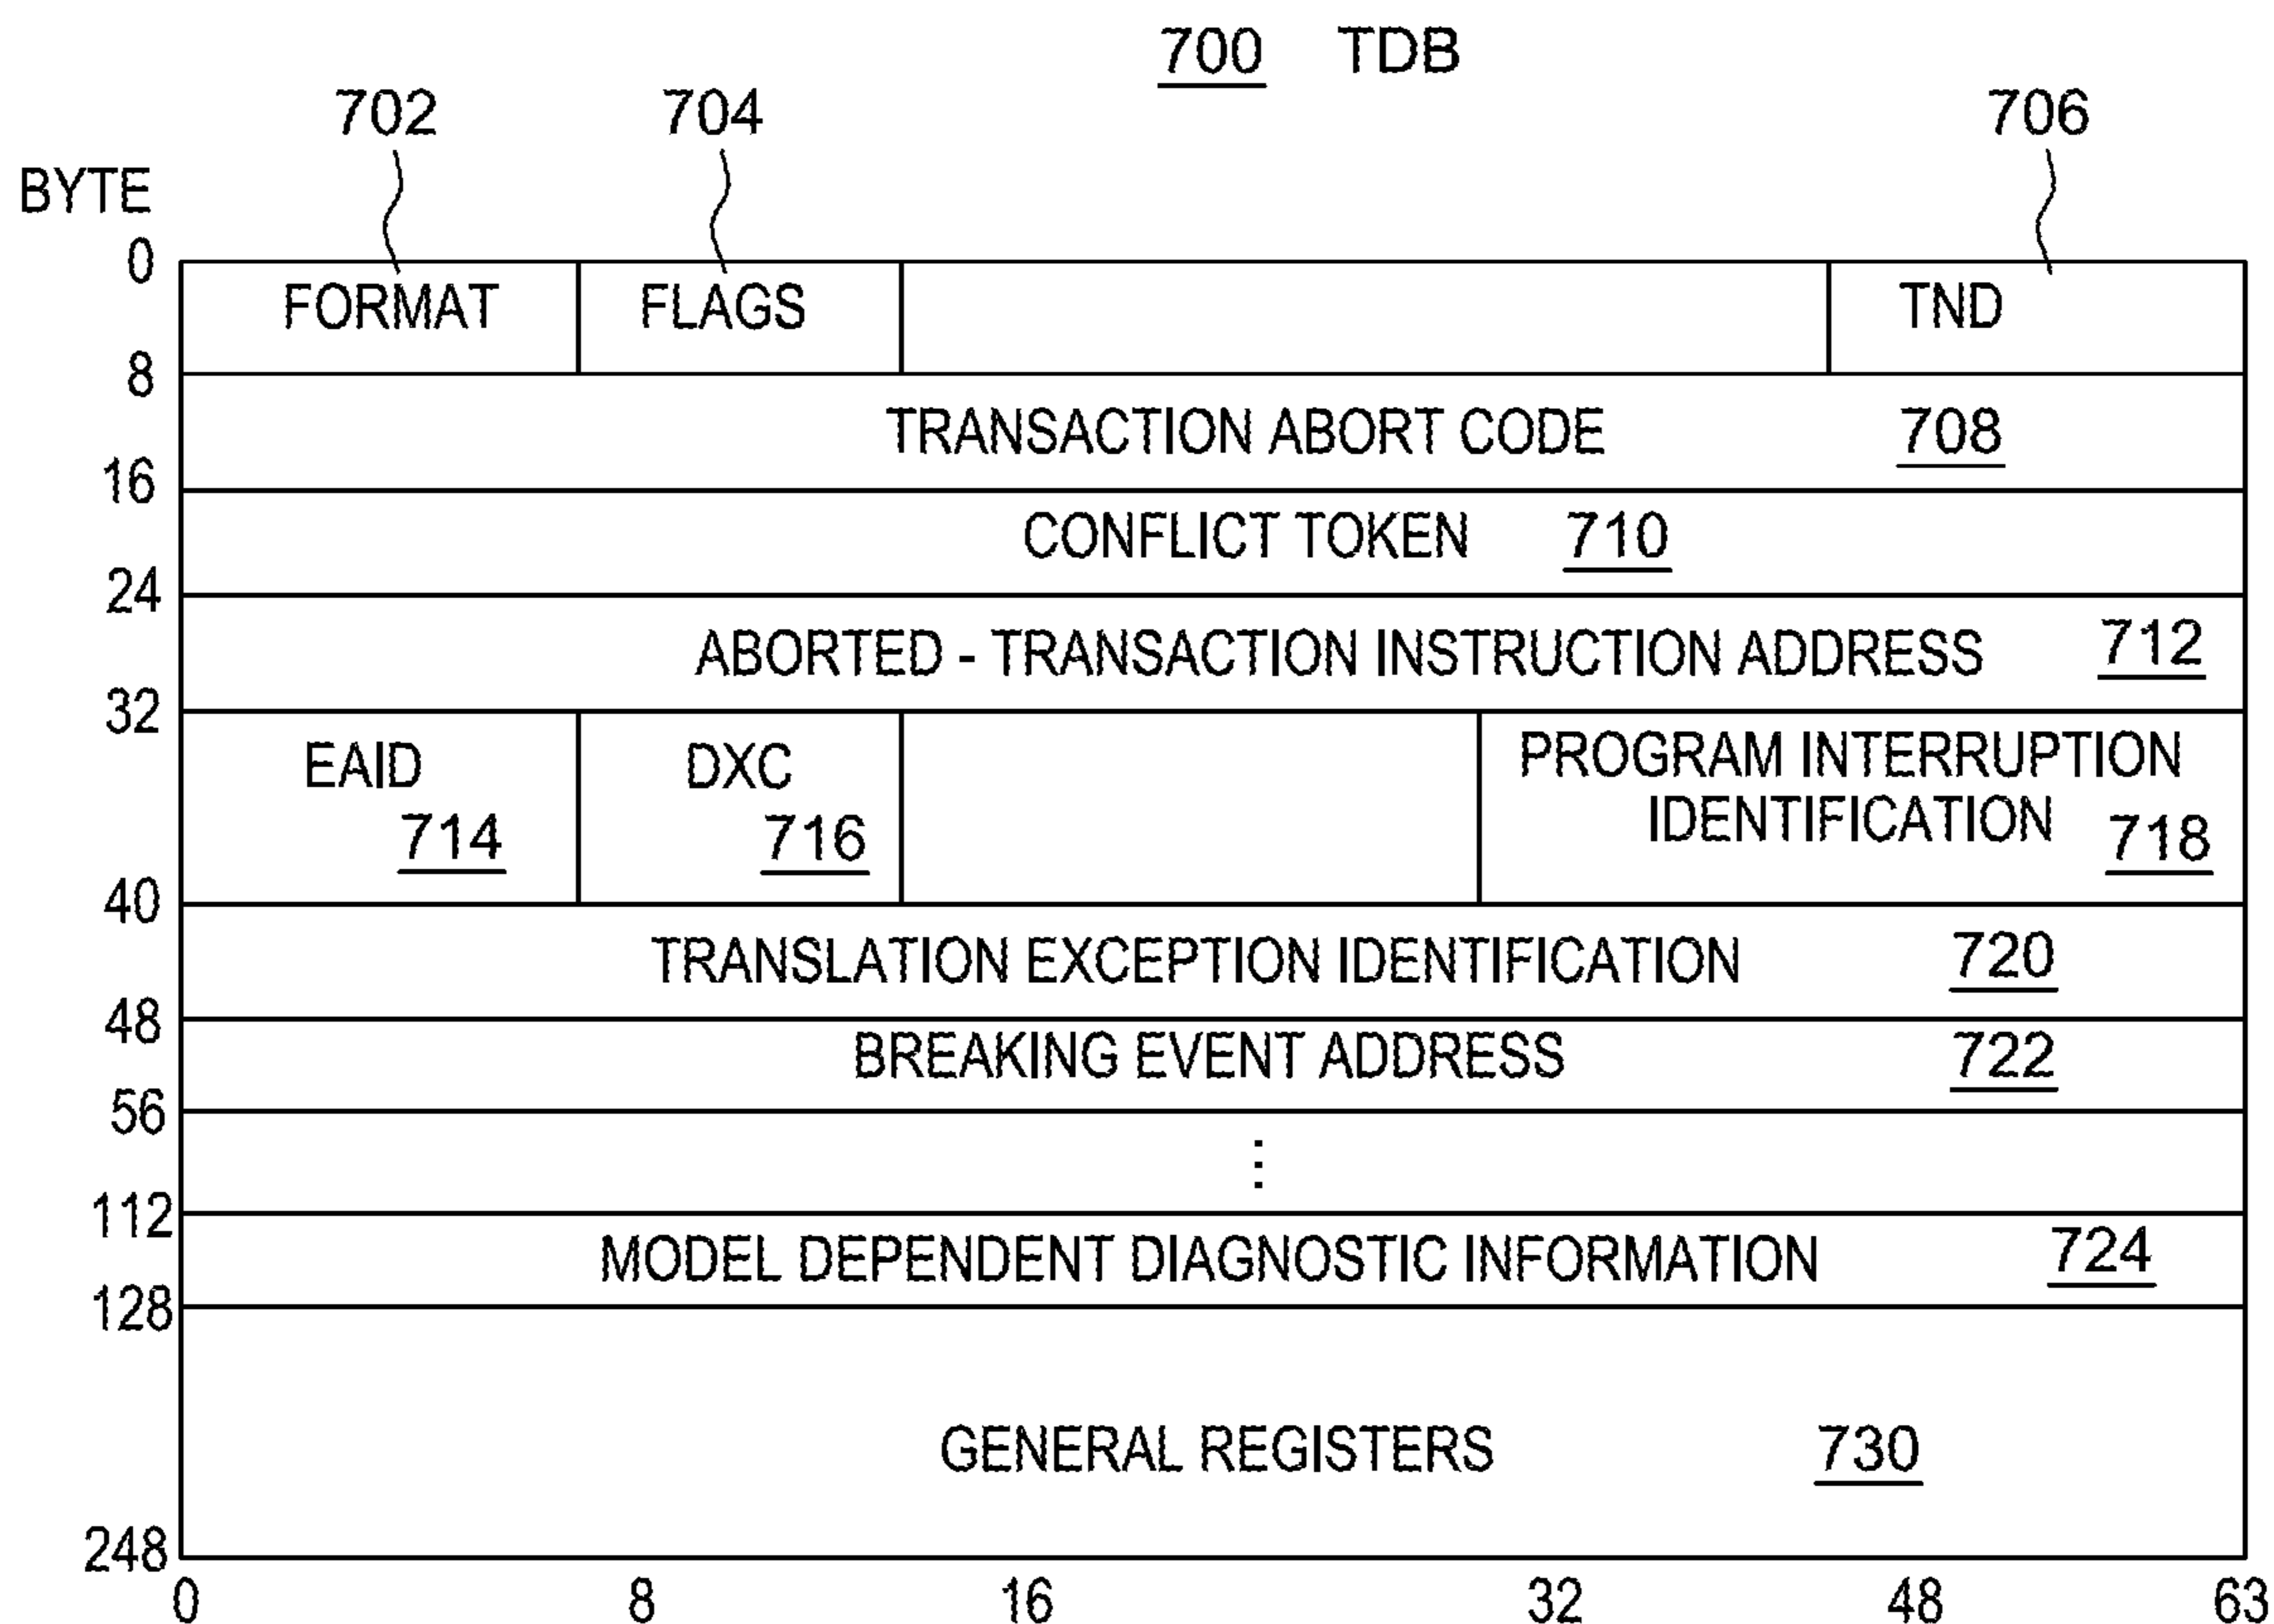


FIG. 7

800

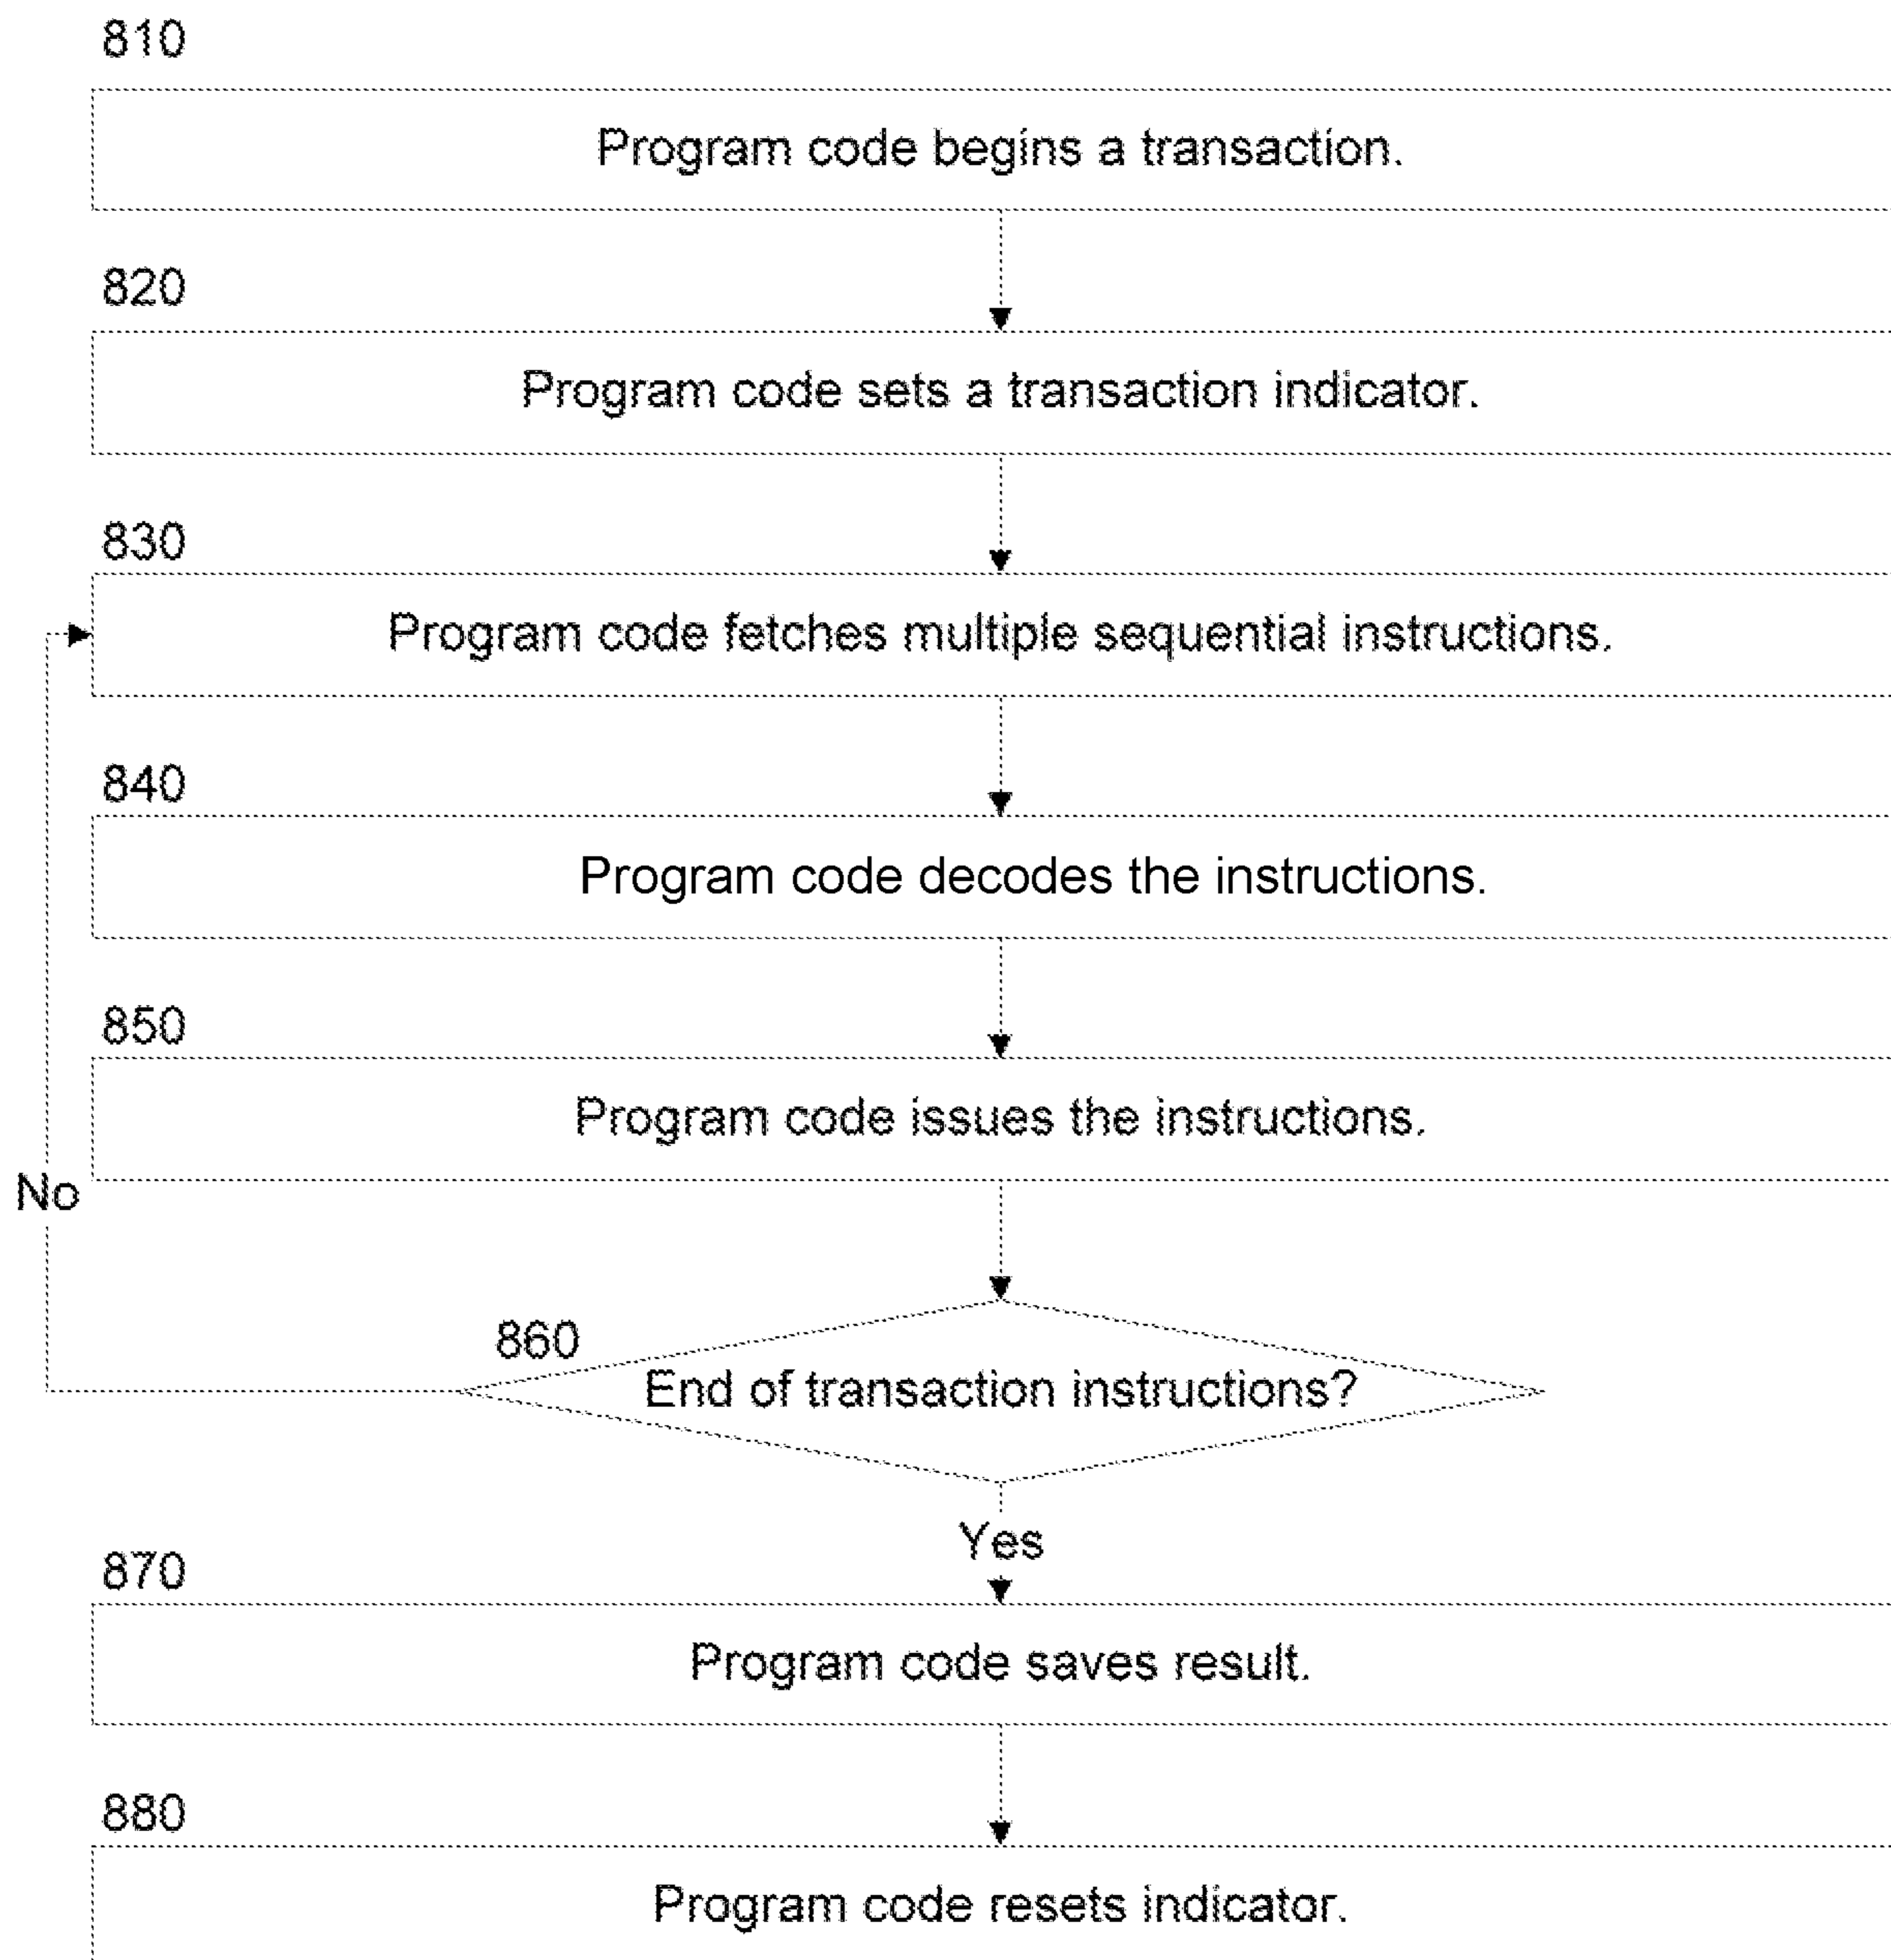


FIG. 8

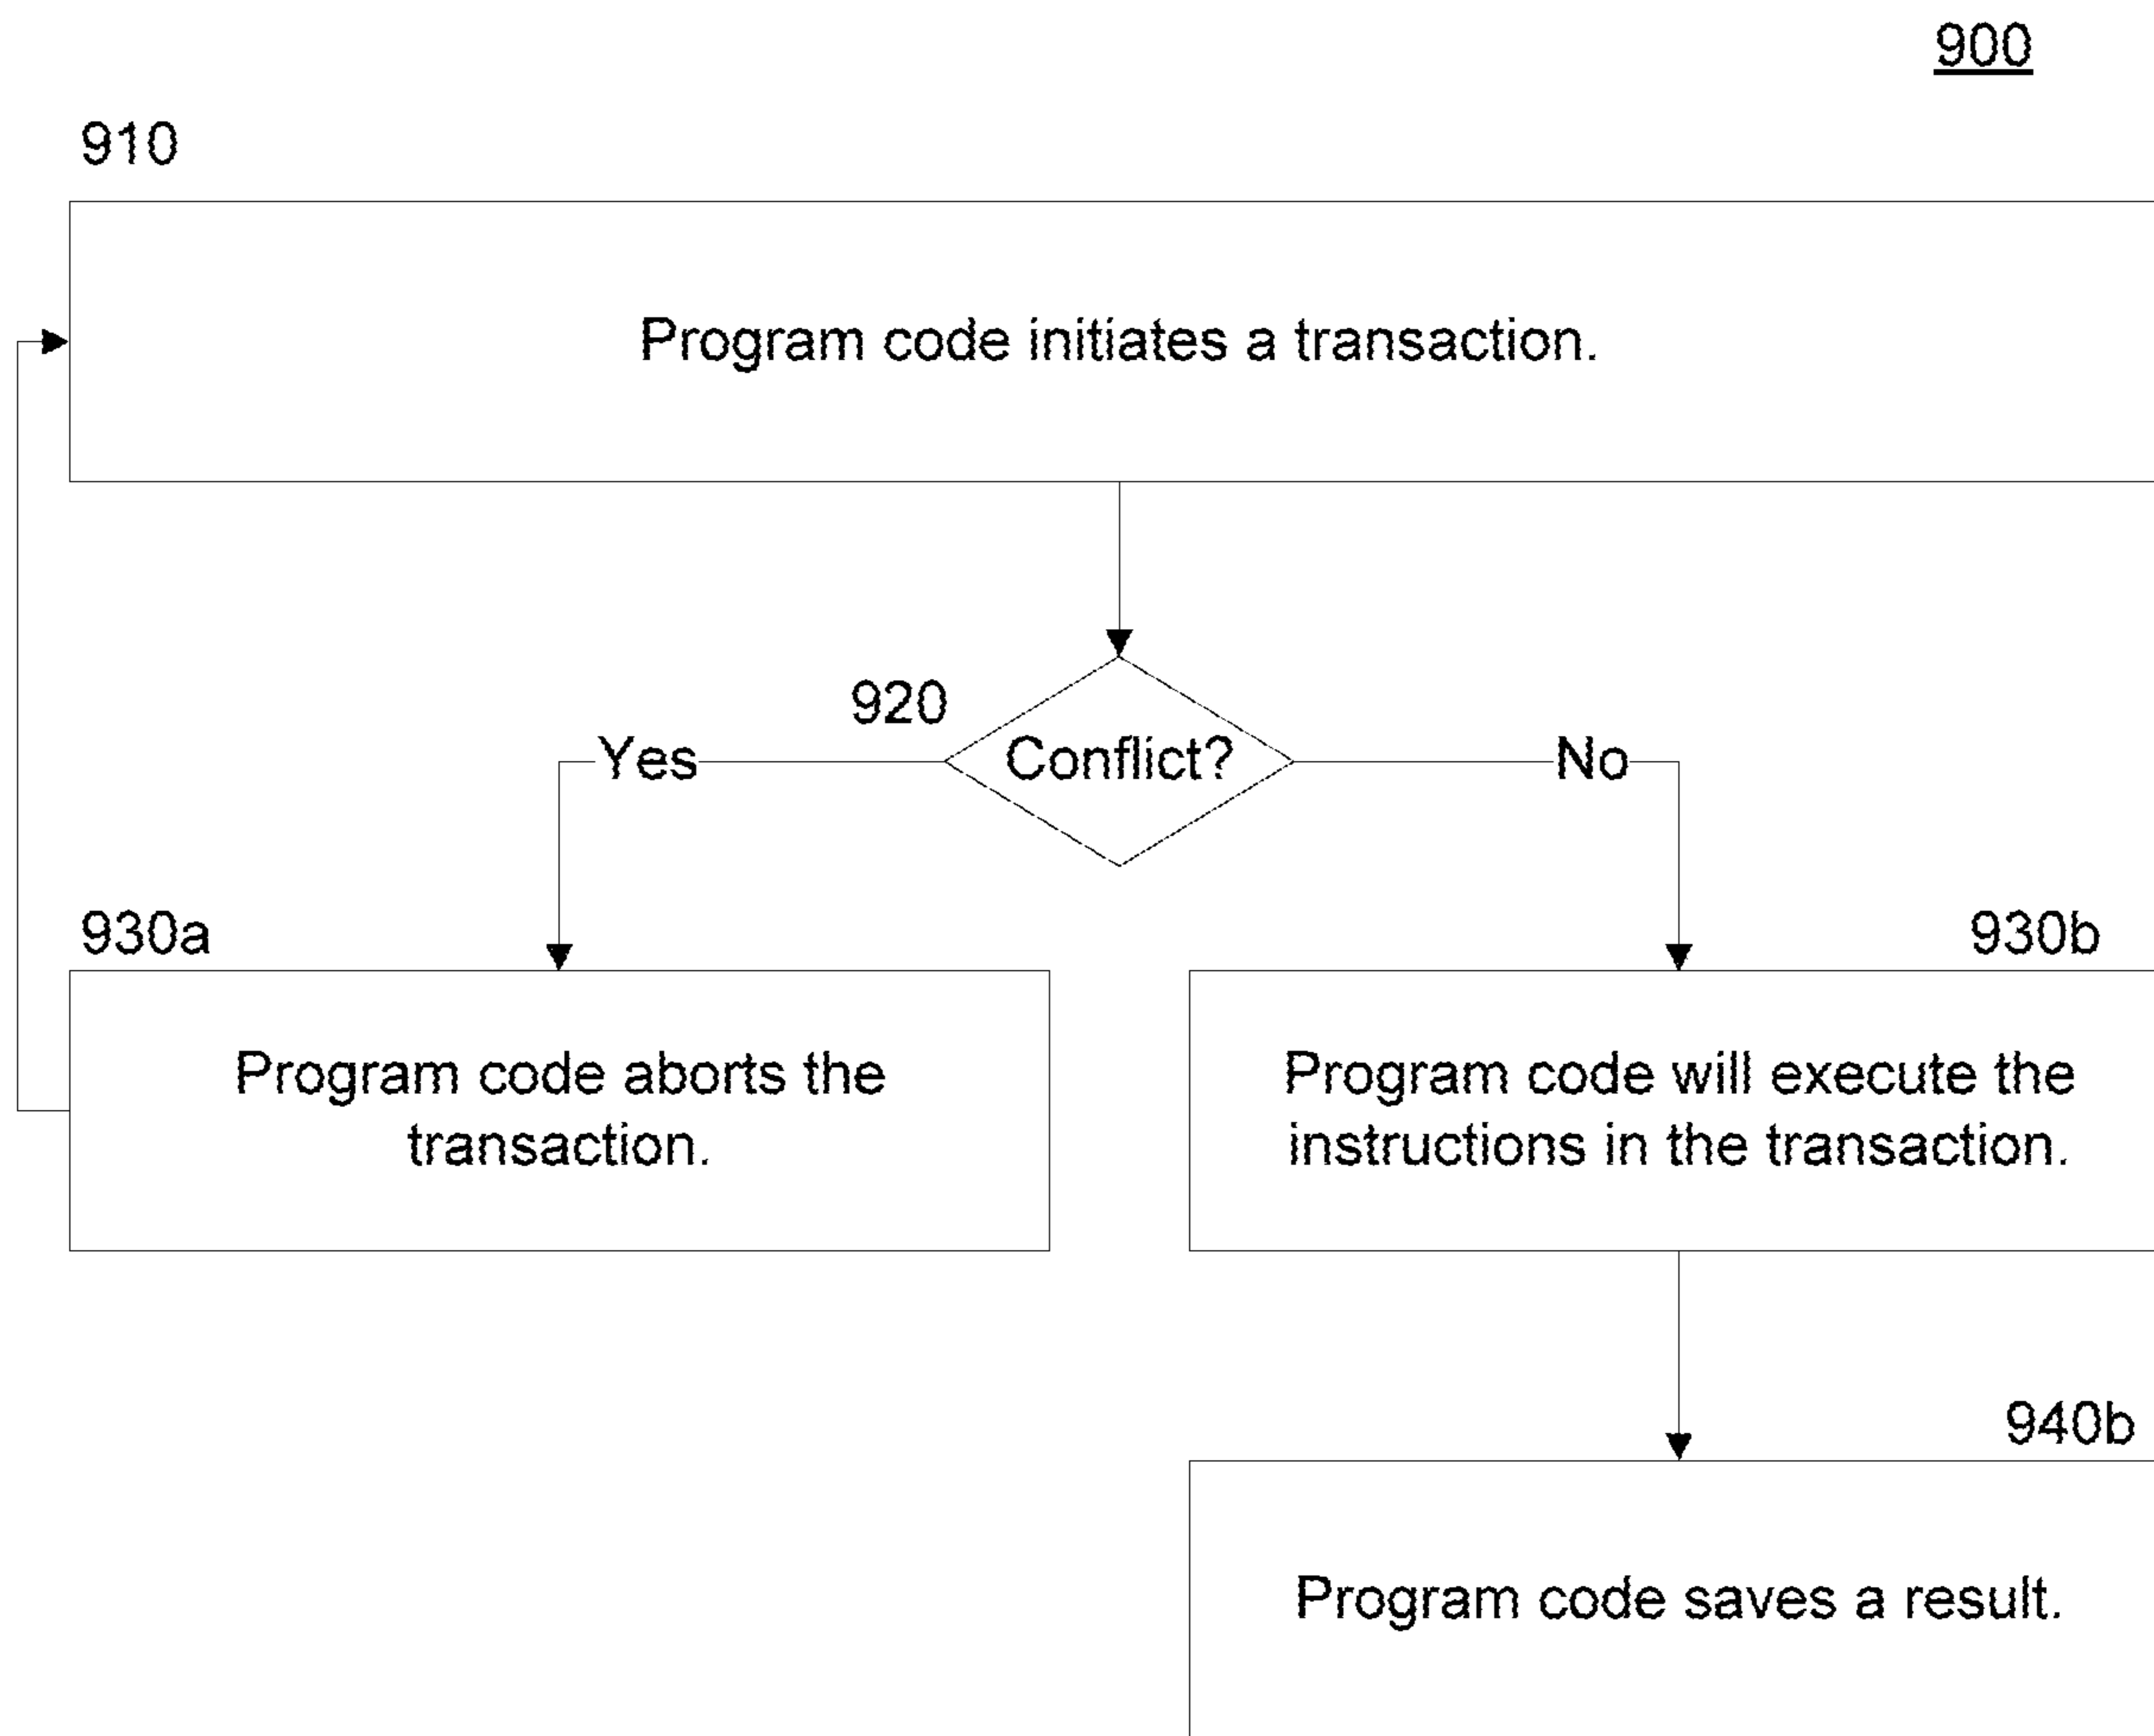


FIG. 9

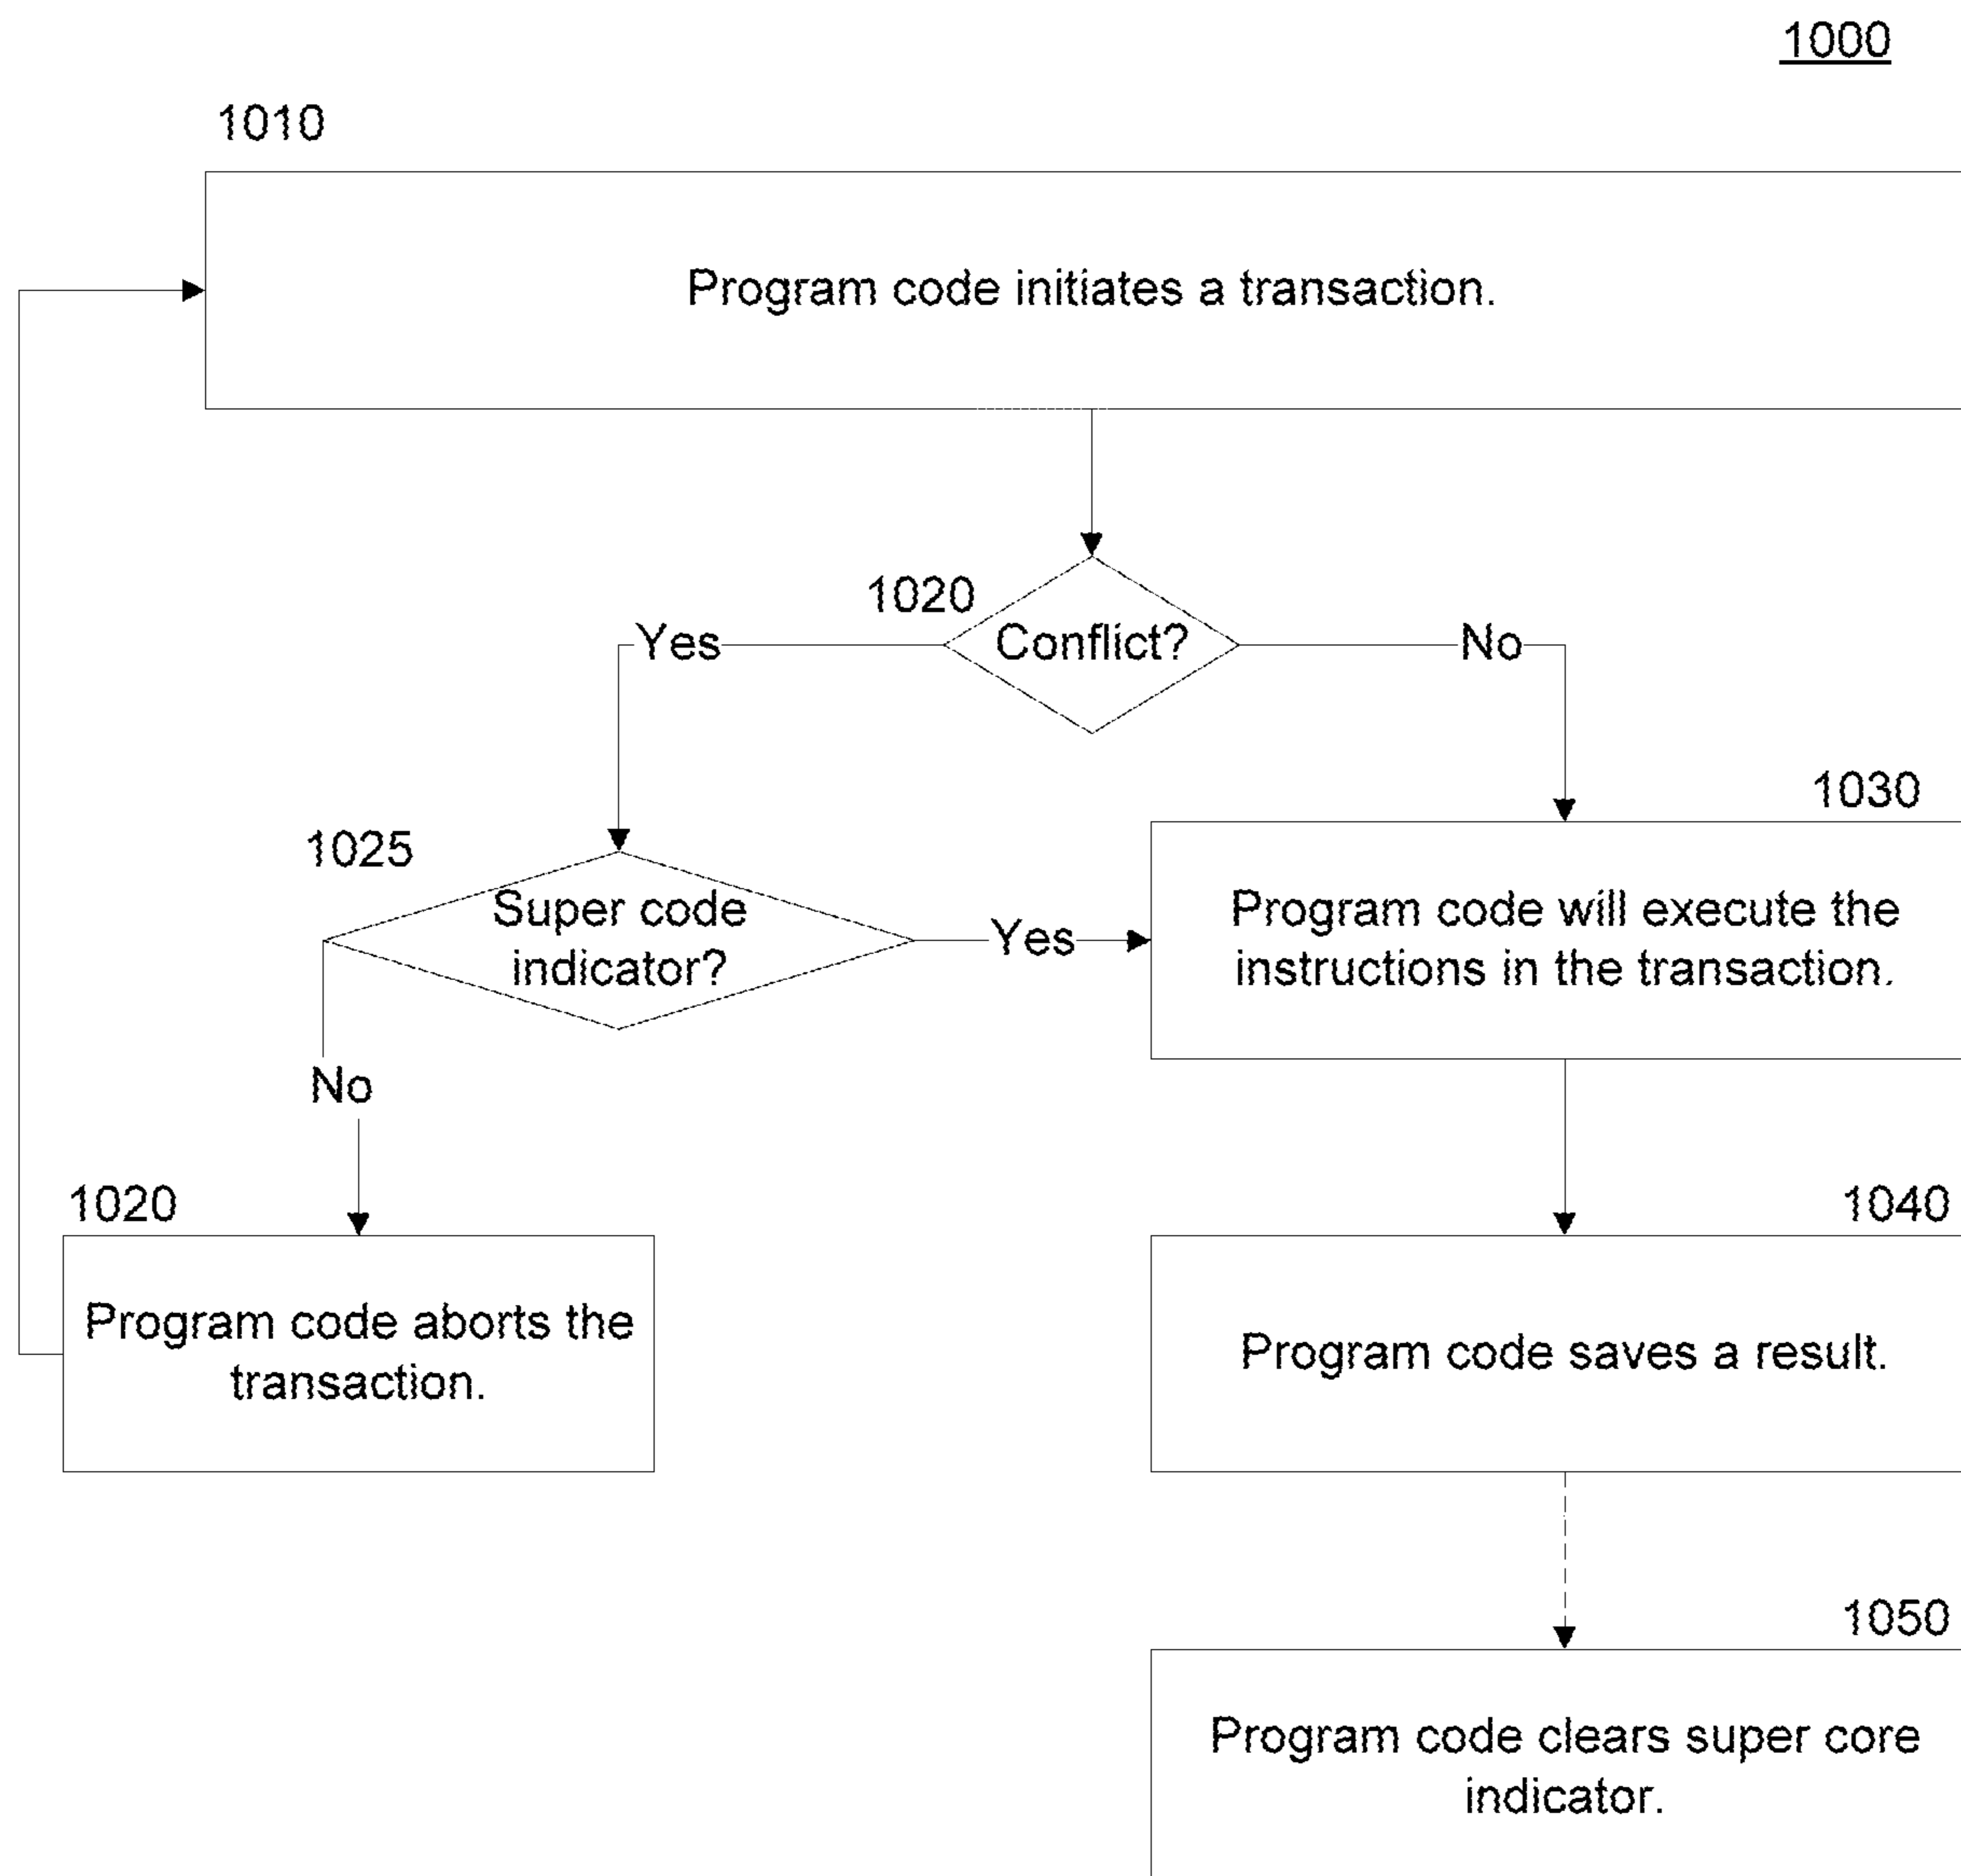


FIG. 10

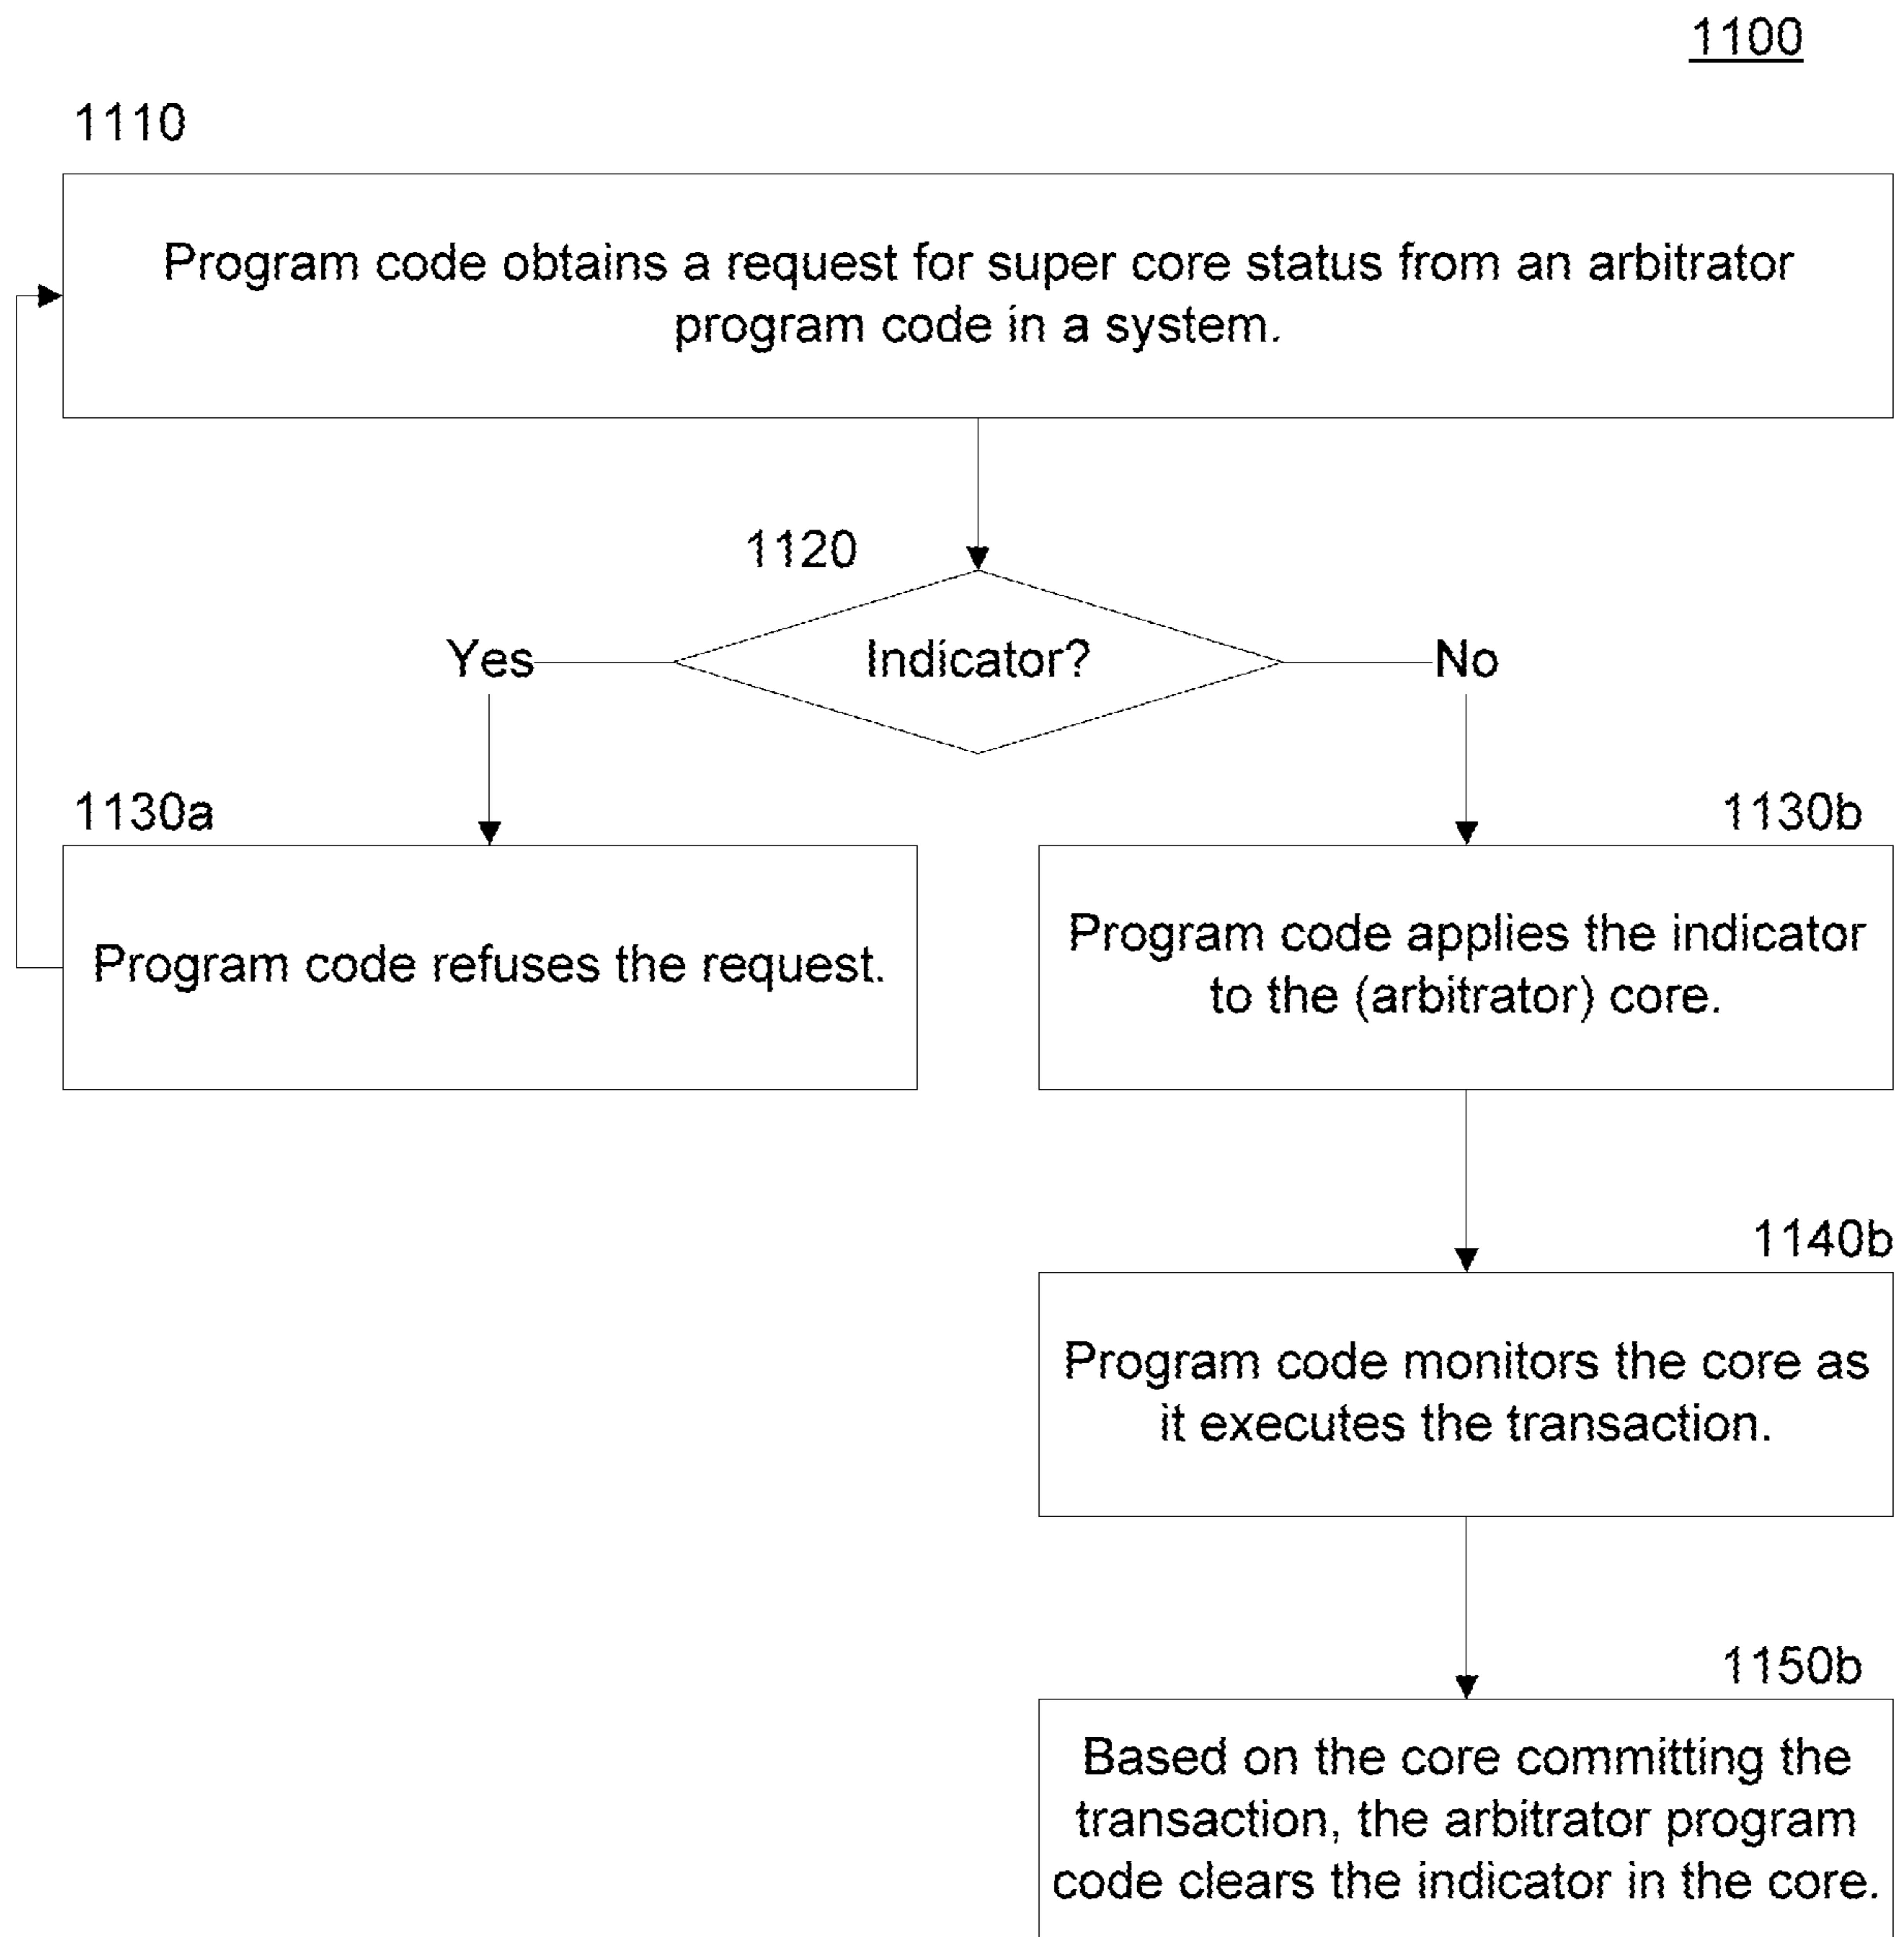


FIG. 11

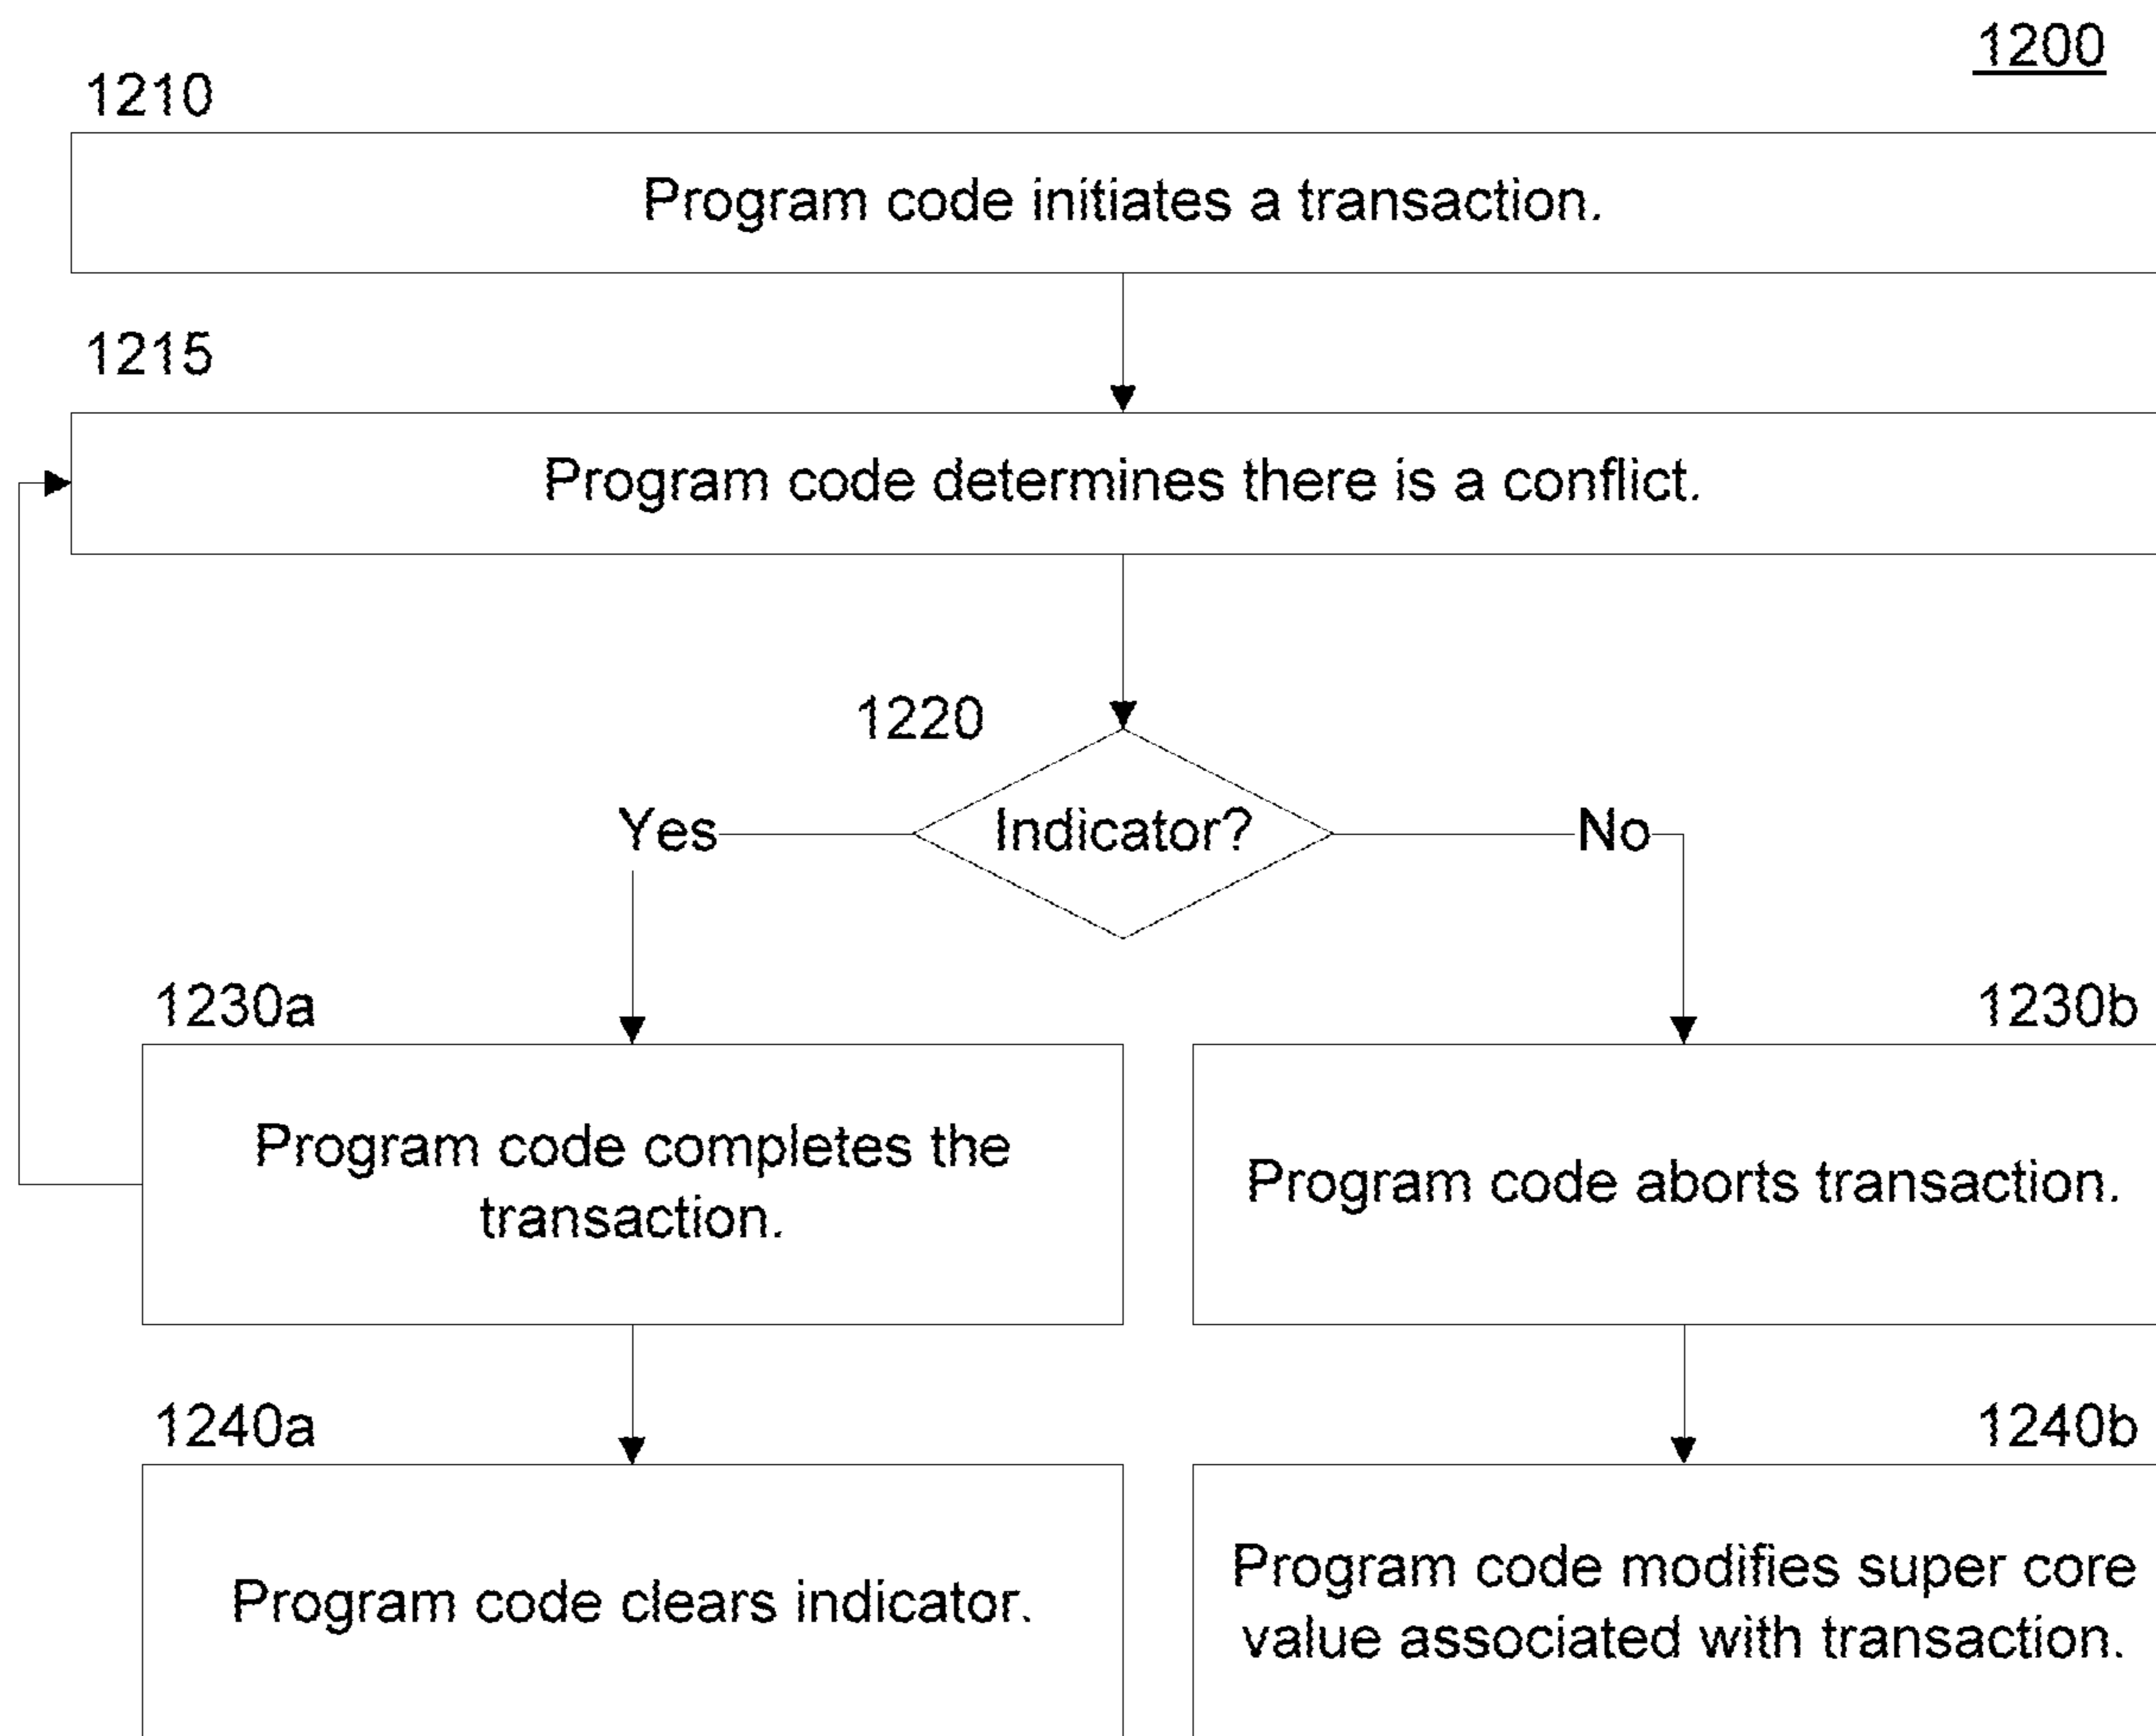


FIG. 12

1300

1310

Program code initiates a transaction and obtains a super core value associated with the transaction to determine whether the super core value is greater than or equal to a threshold value.

1320

Based on determining that the super core value is greater than or equal to the threshold, the program code requests super core status for the core executing the transaction.

1330

The program code obtains the super core status (once available) and the program code sets a super core indicator in the core.

FIG. 13

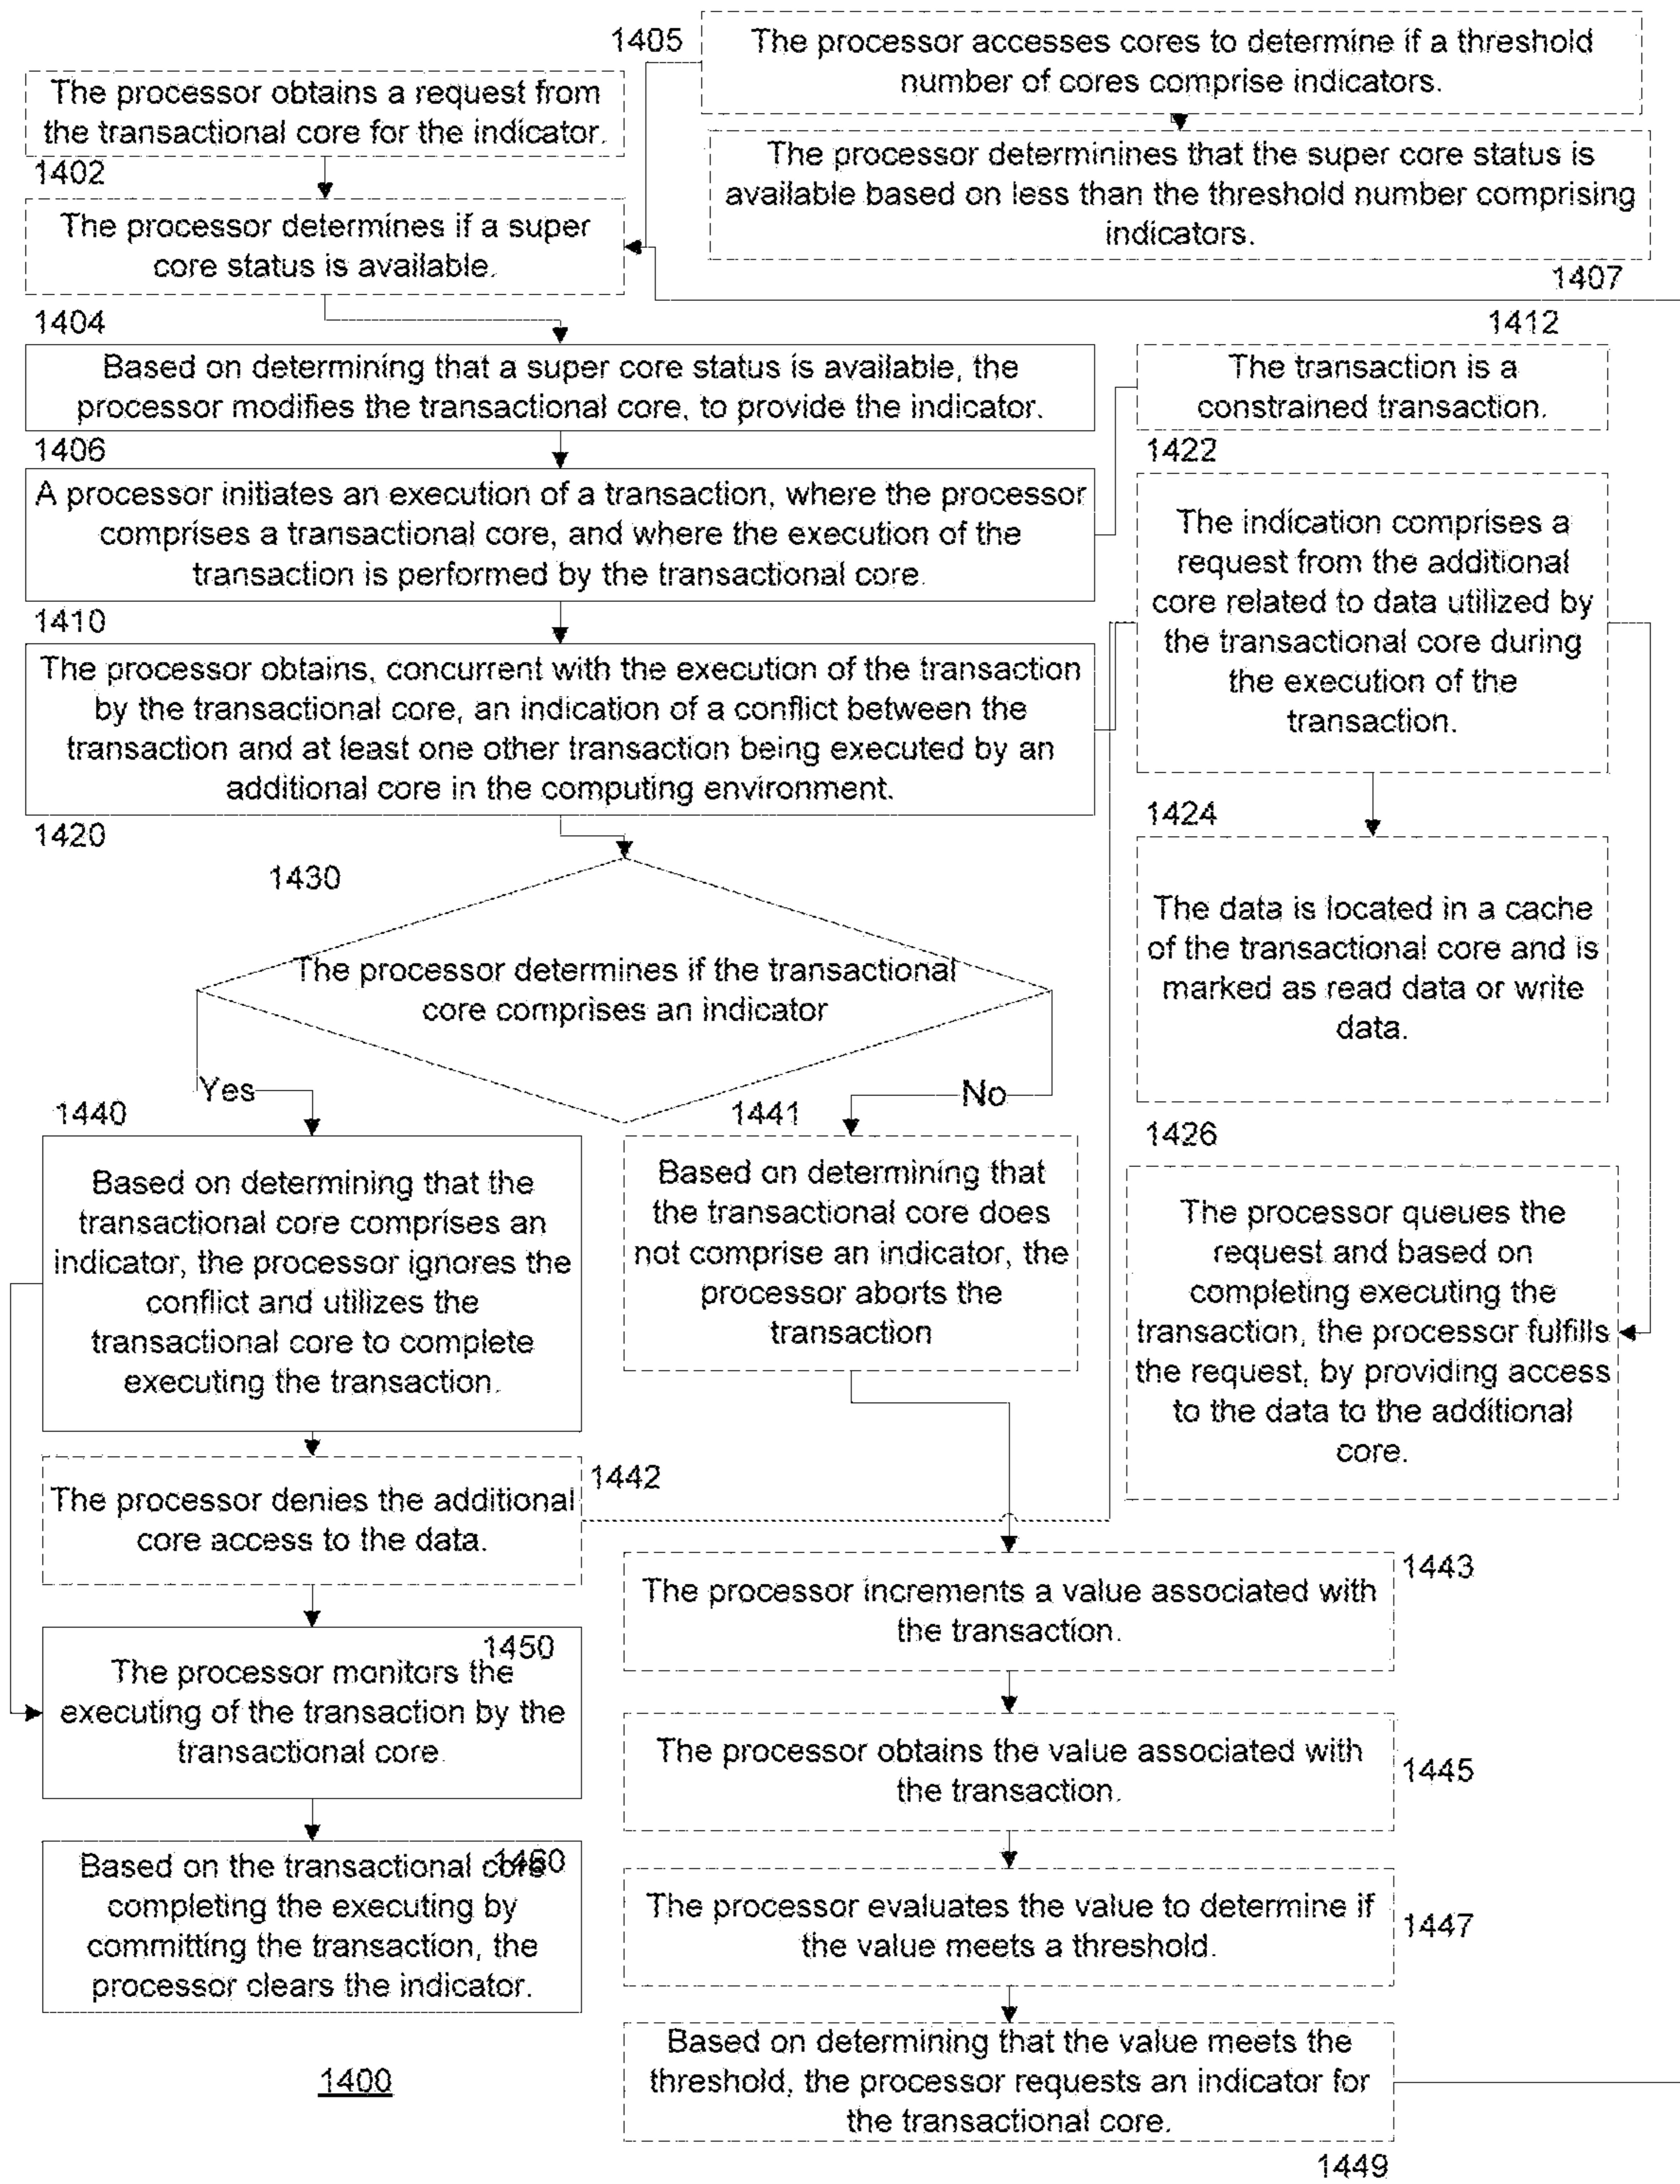


FIG. 14

COMPUTER
PROGRAM
PRODUCT
1200

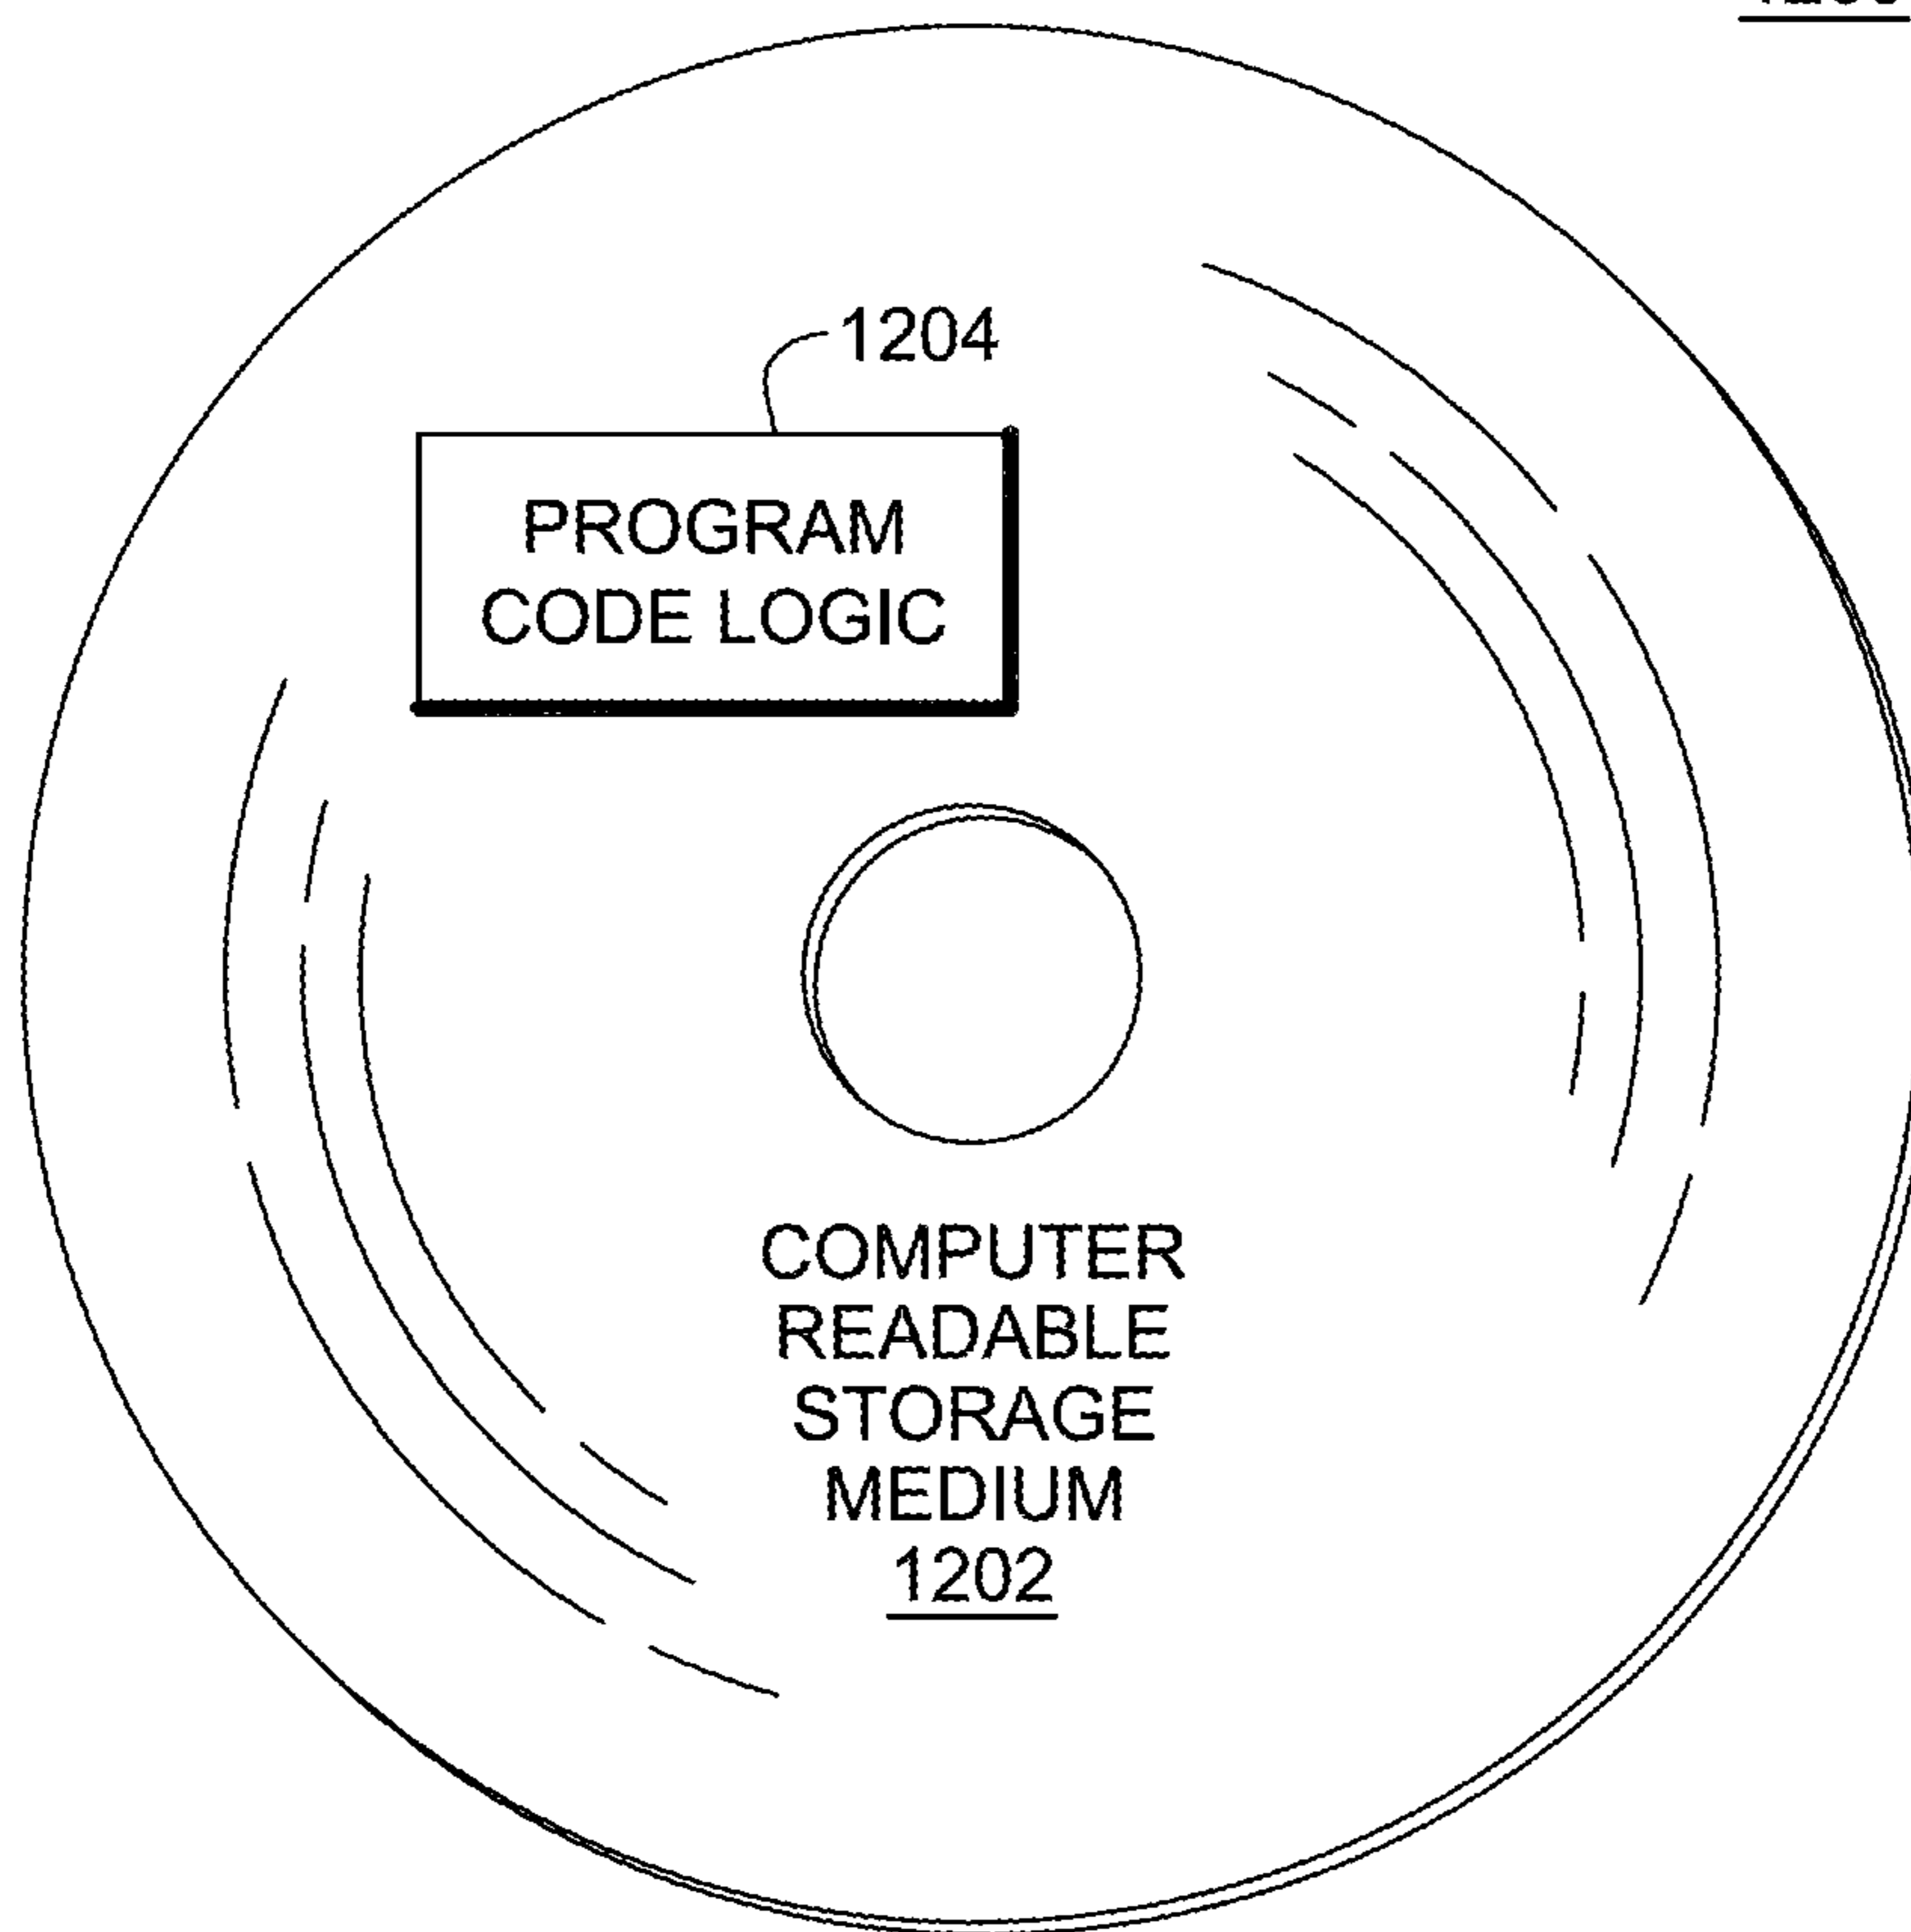


FIG. 15

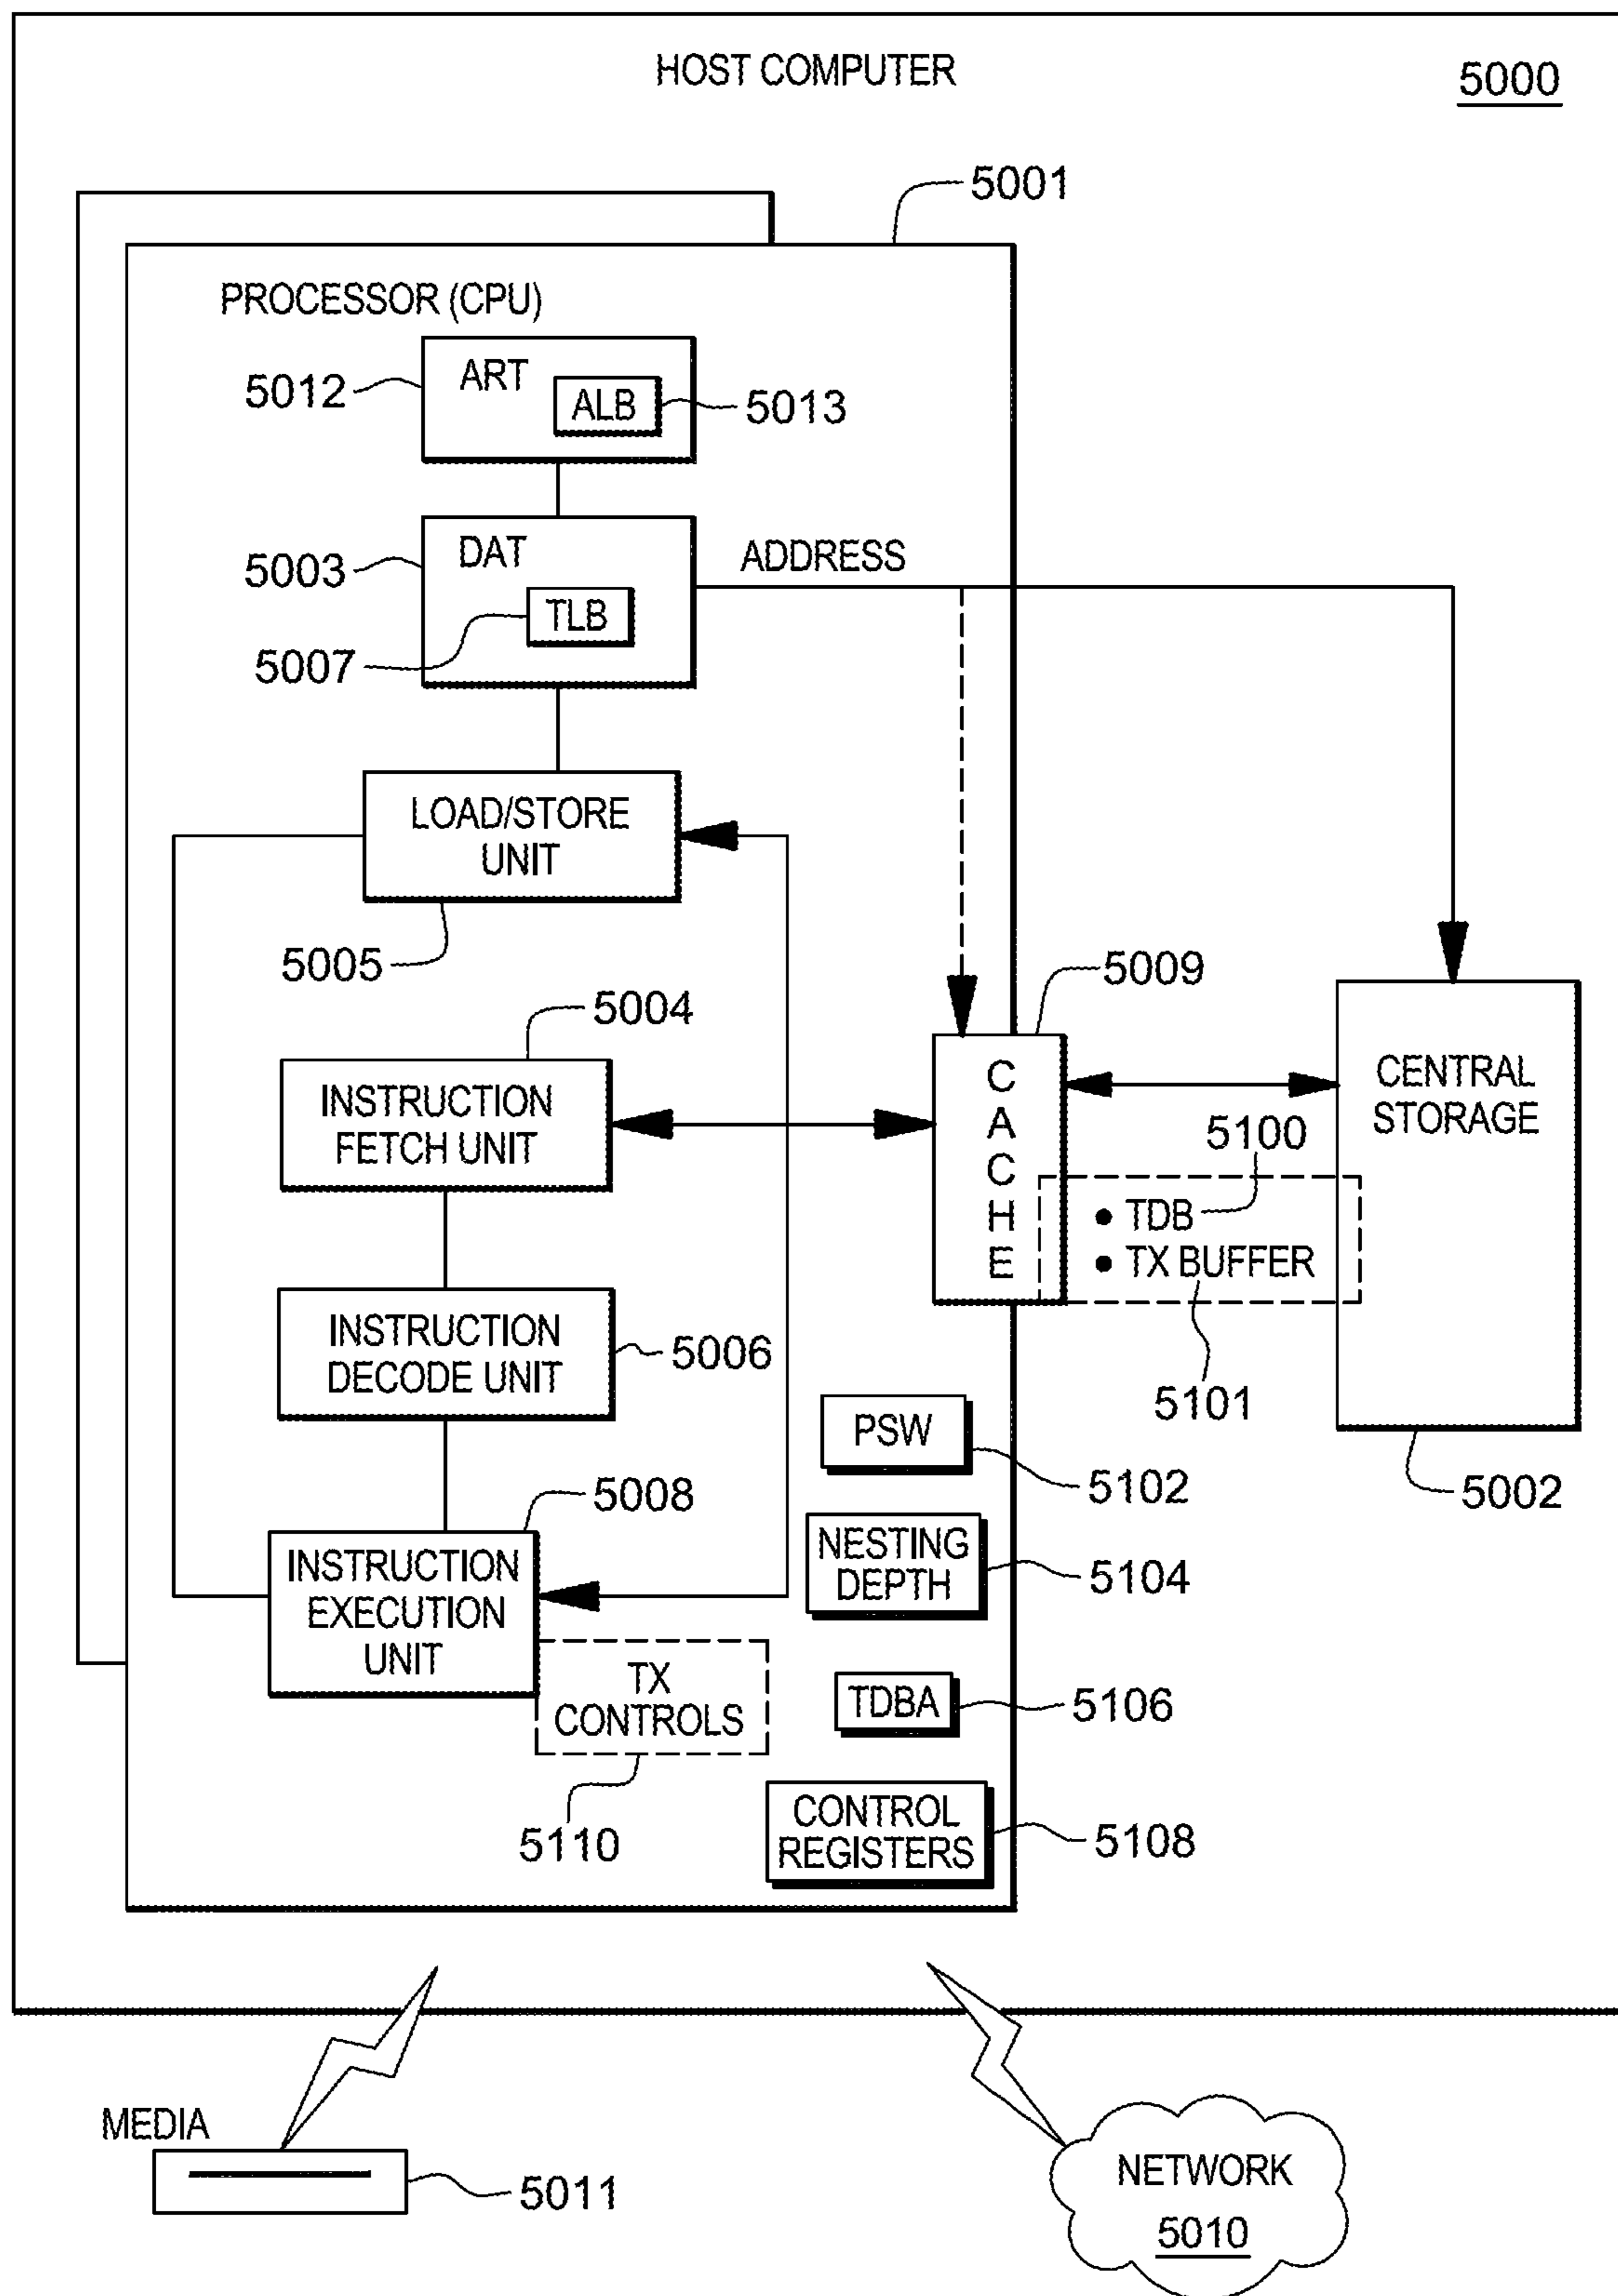


FIG. 16

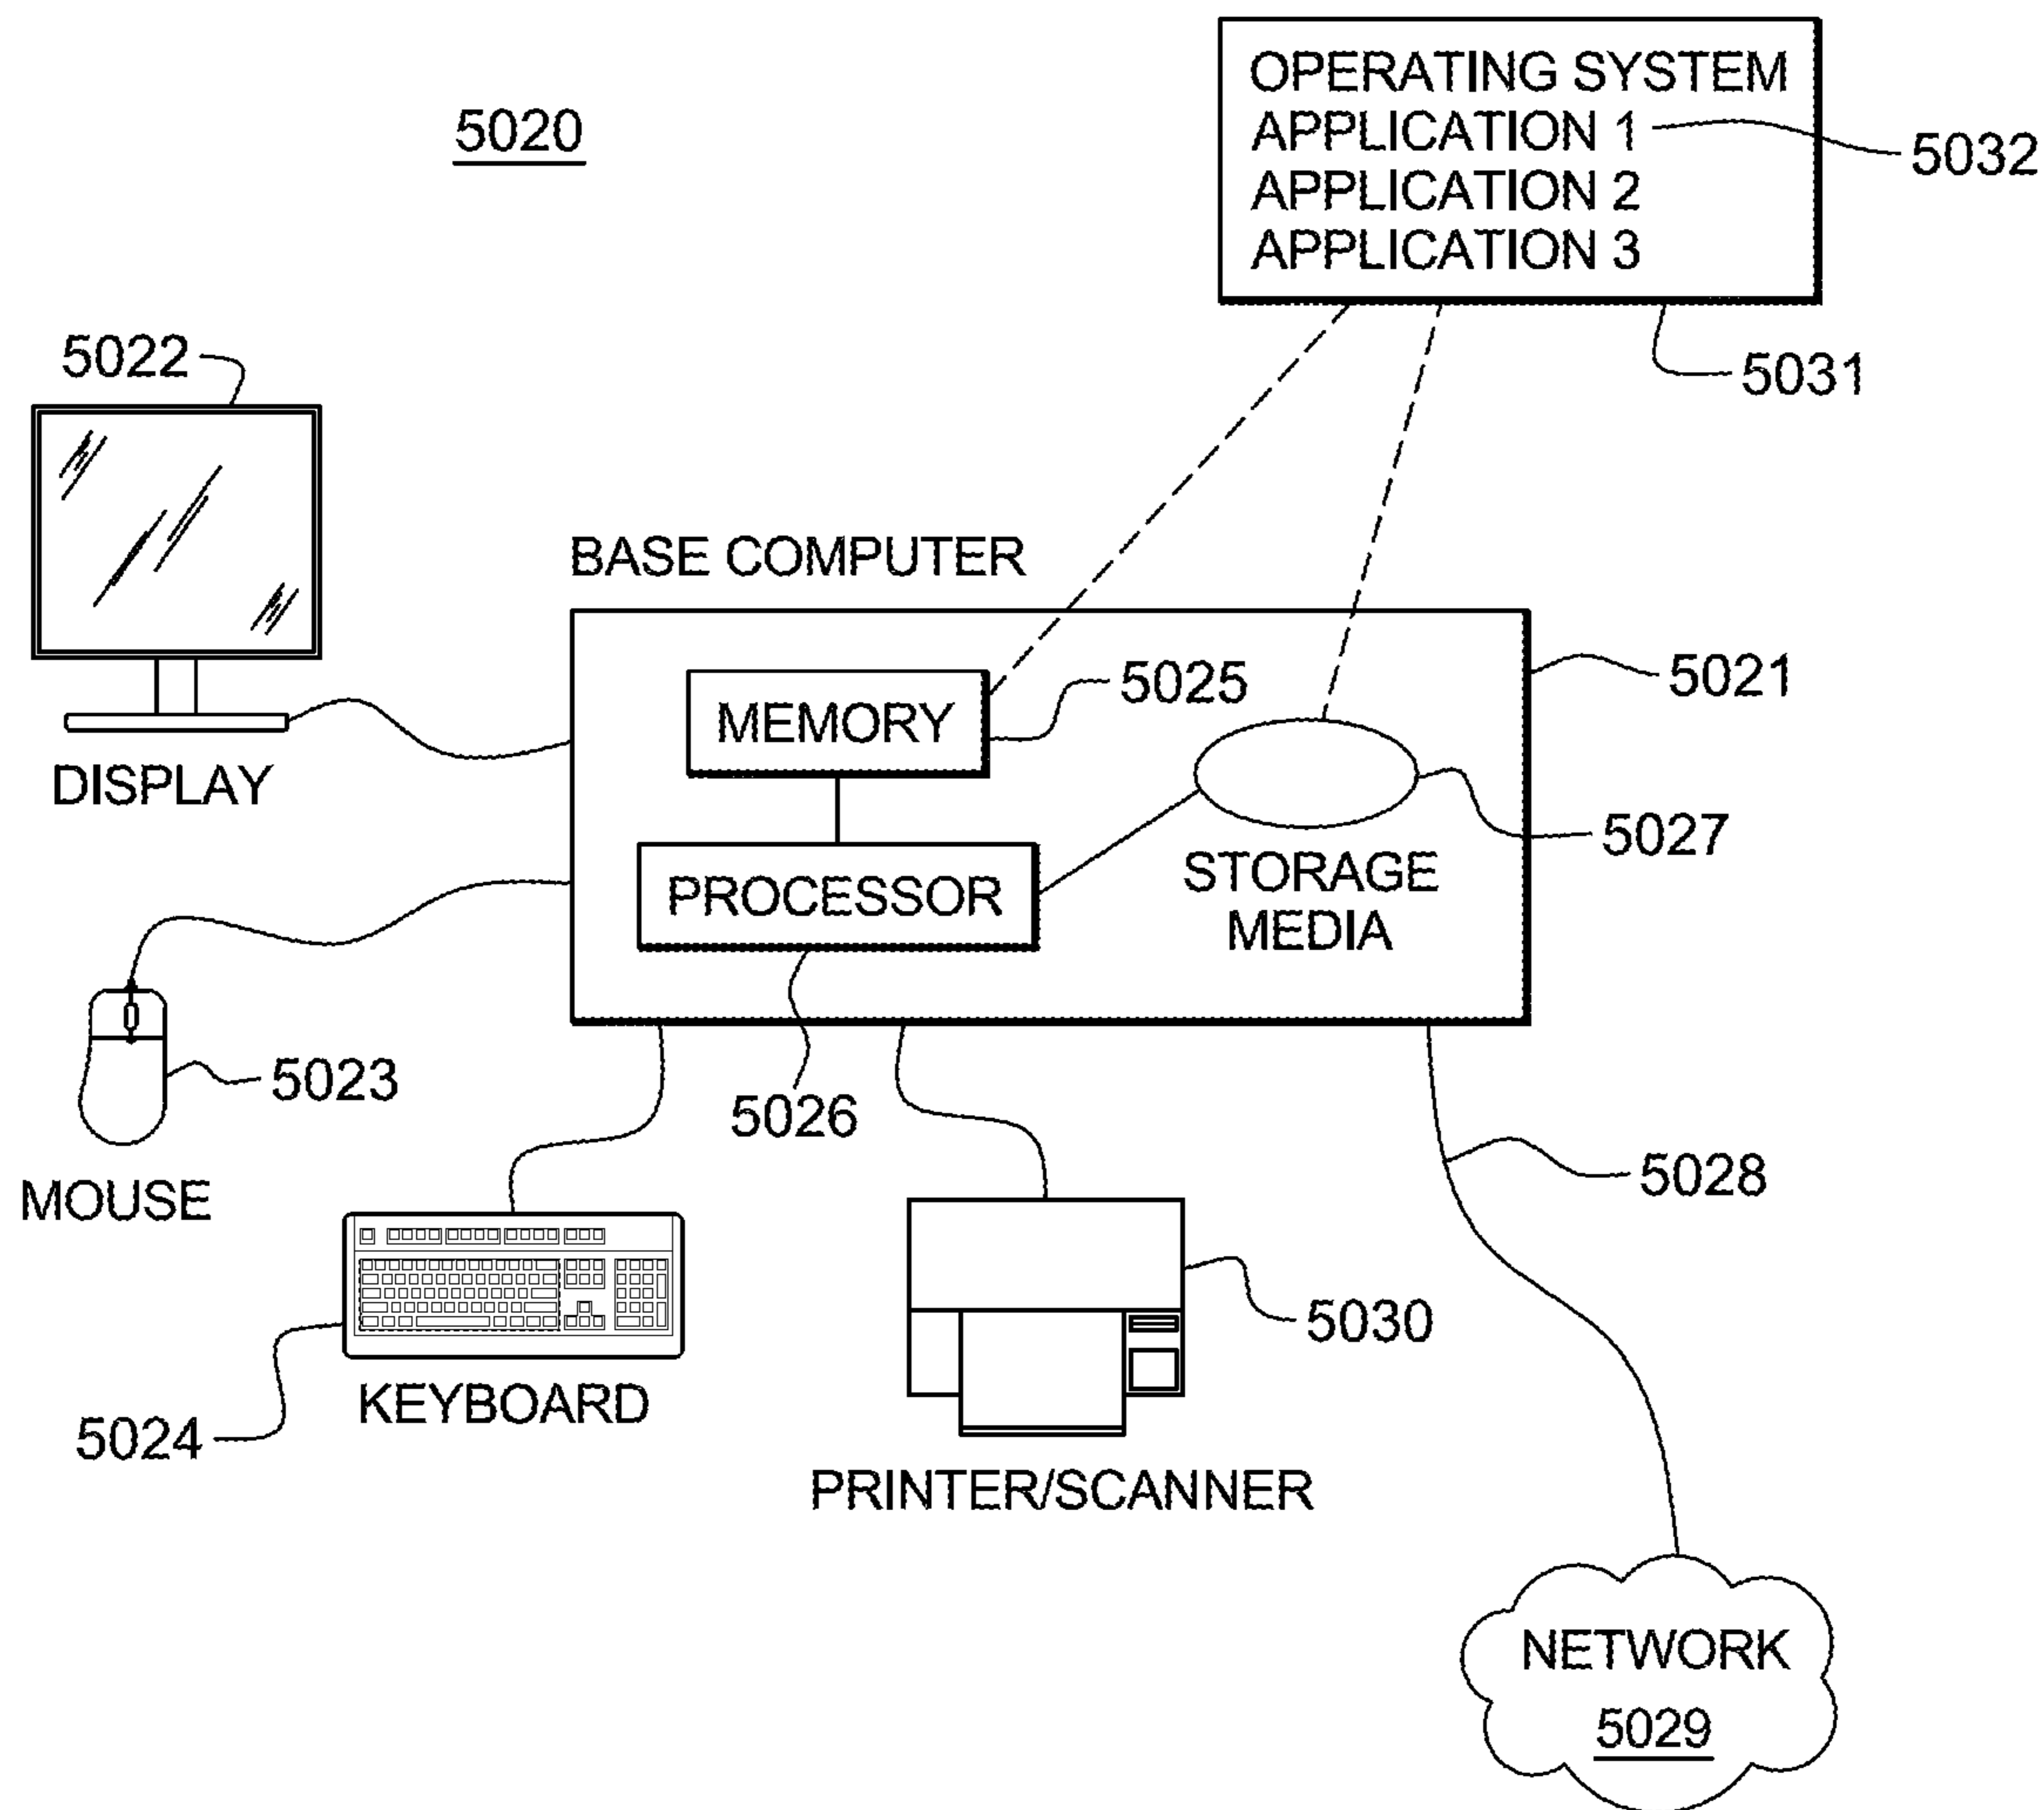


FIG. 17

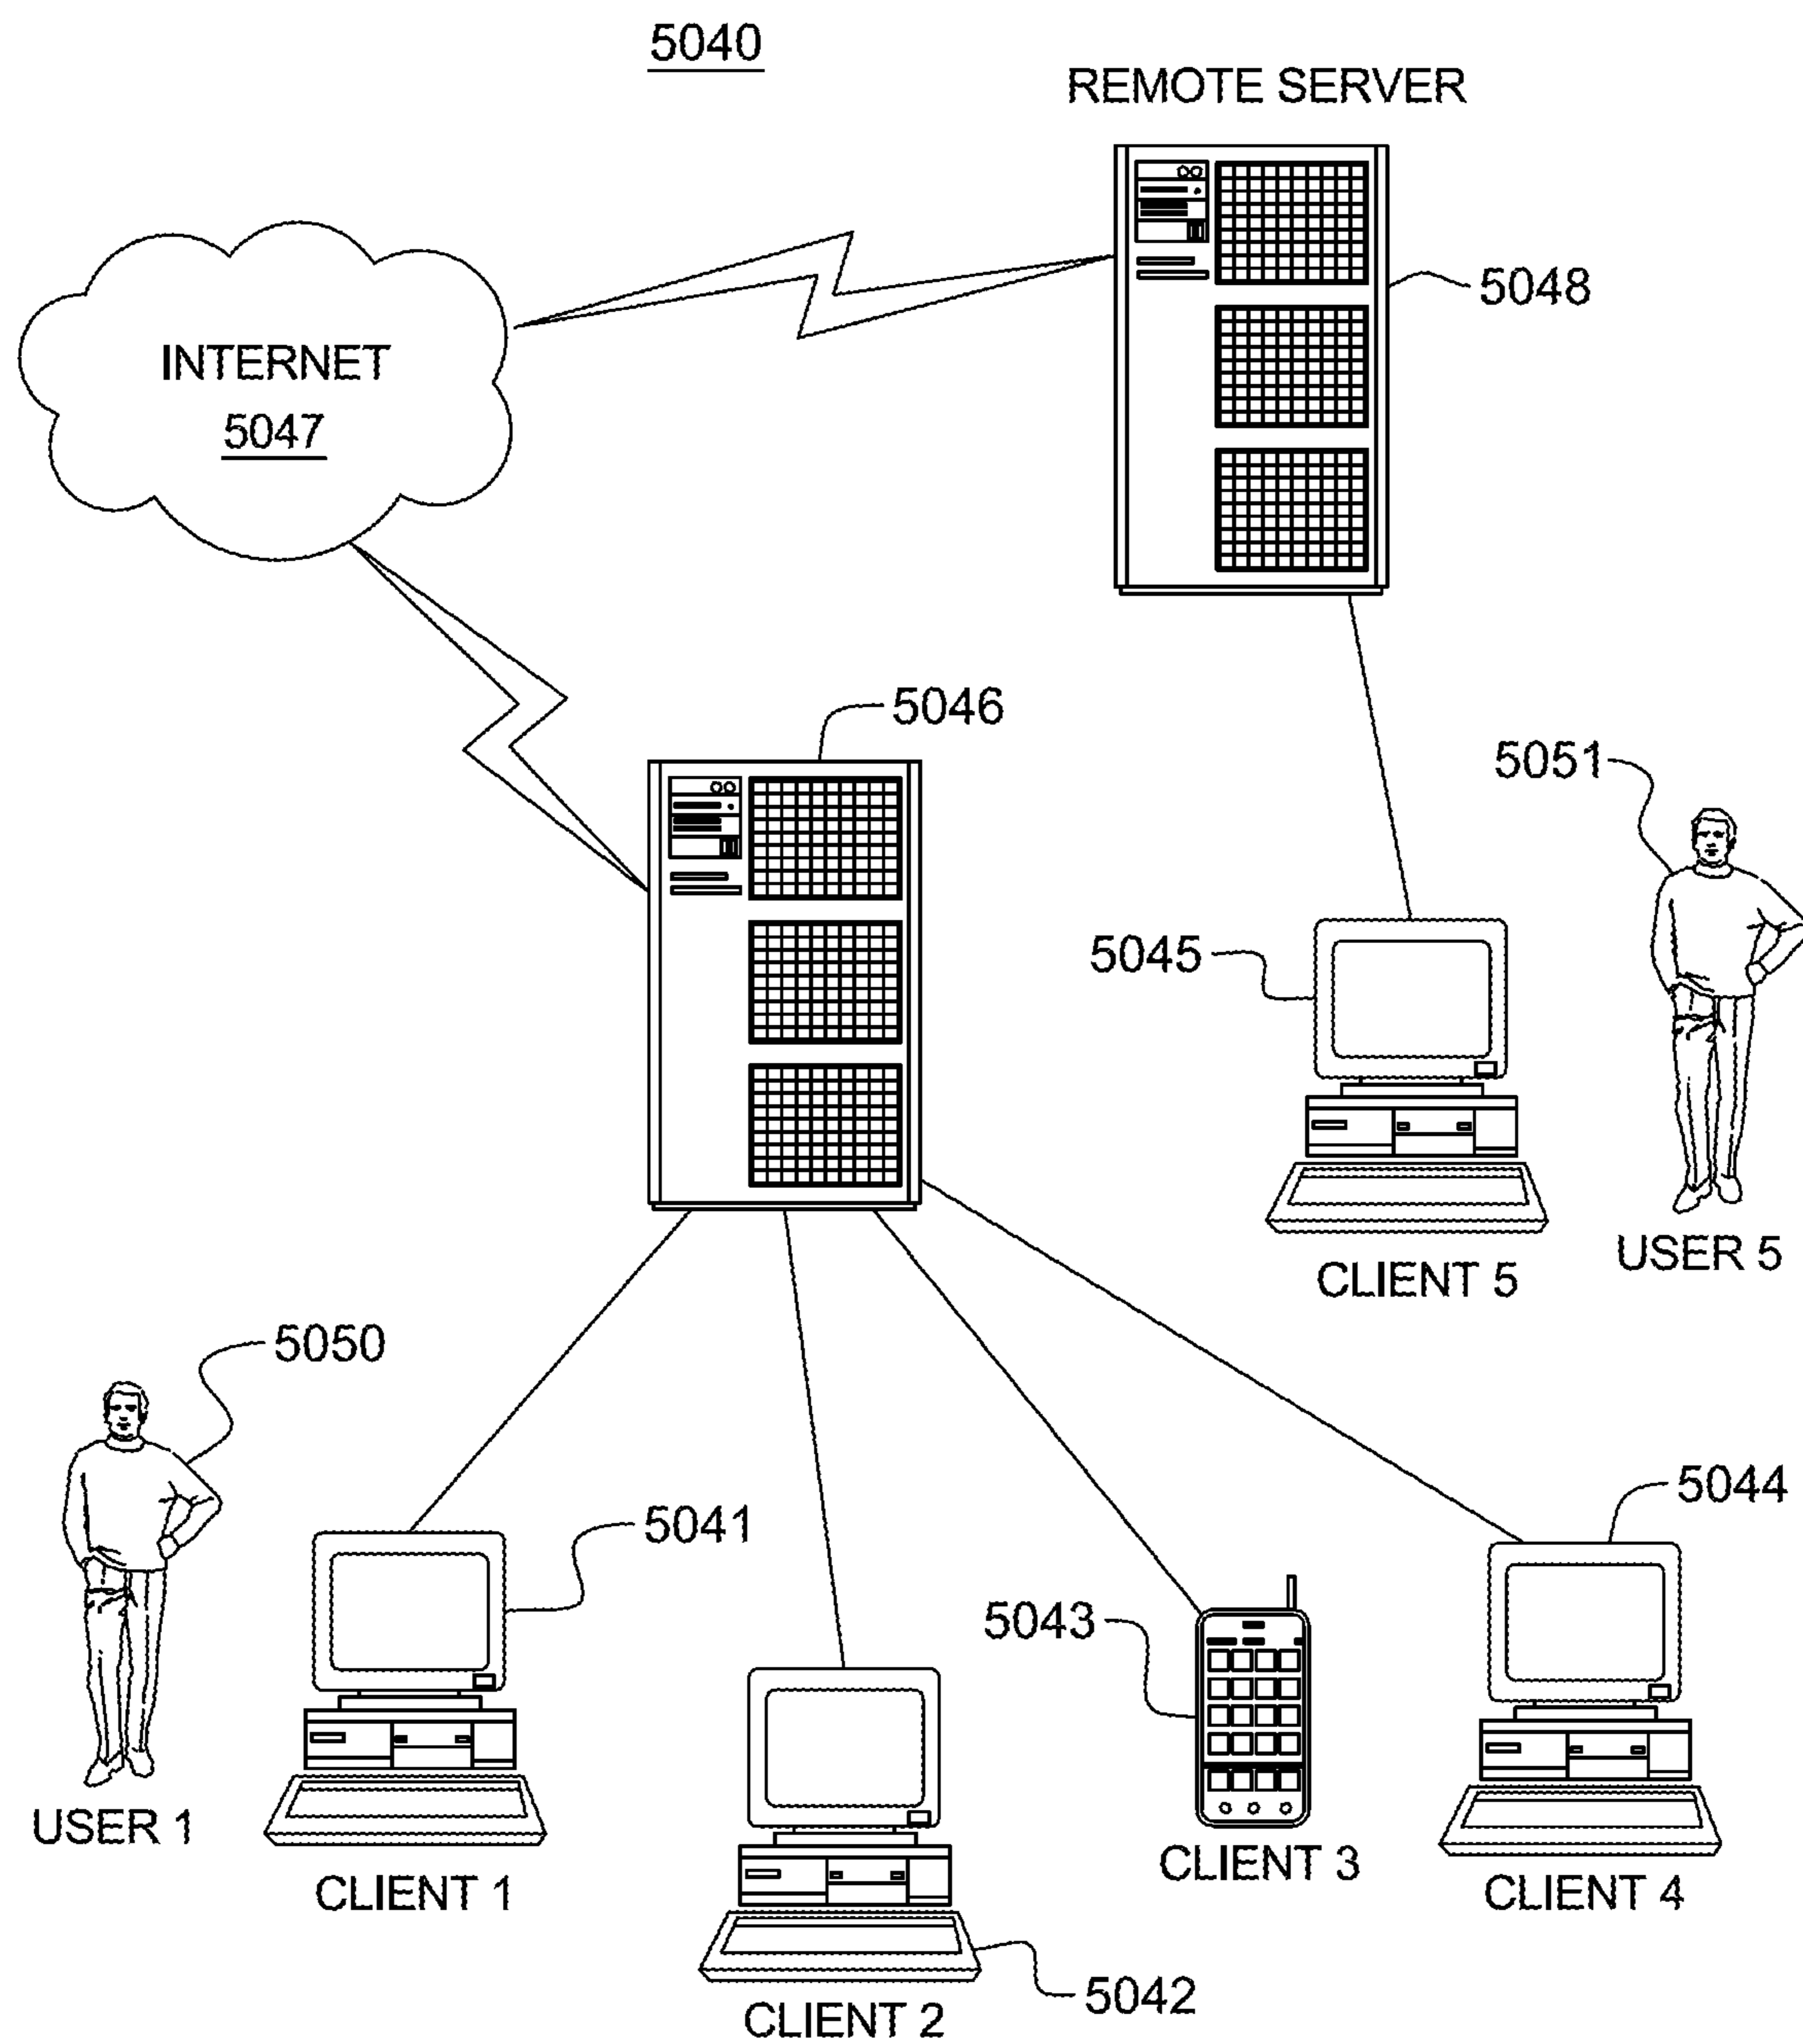


FIG. 18

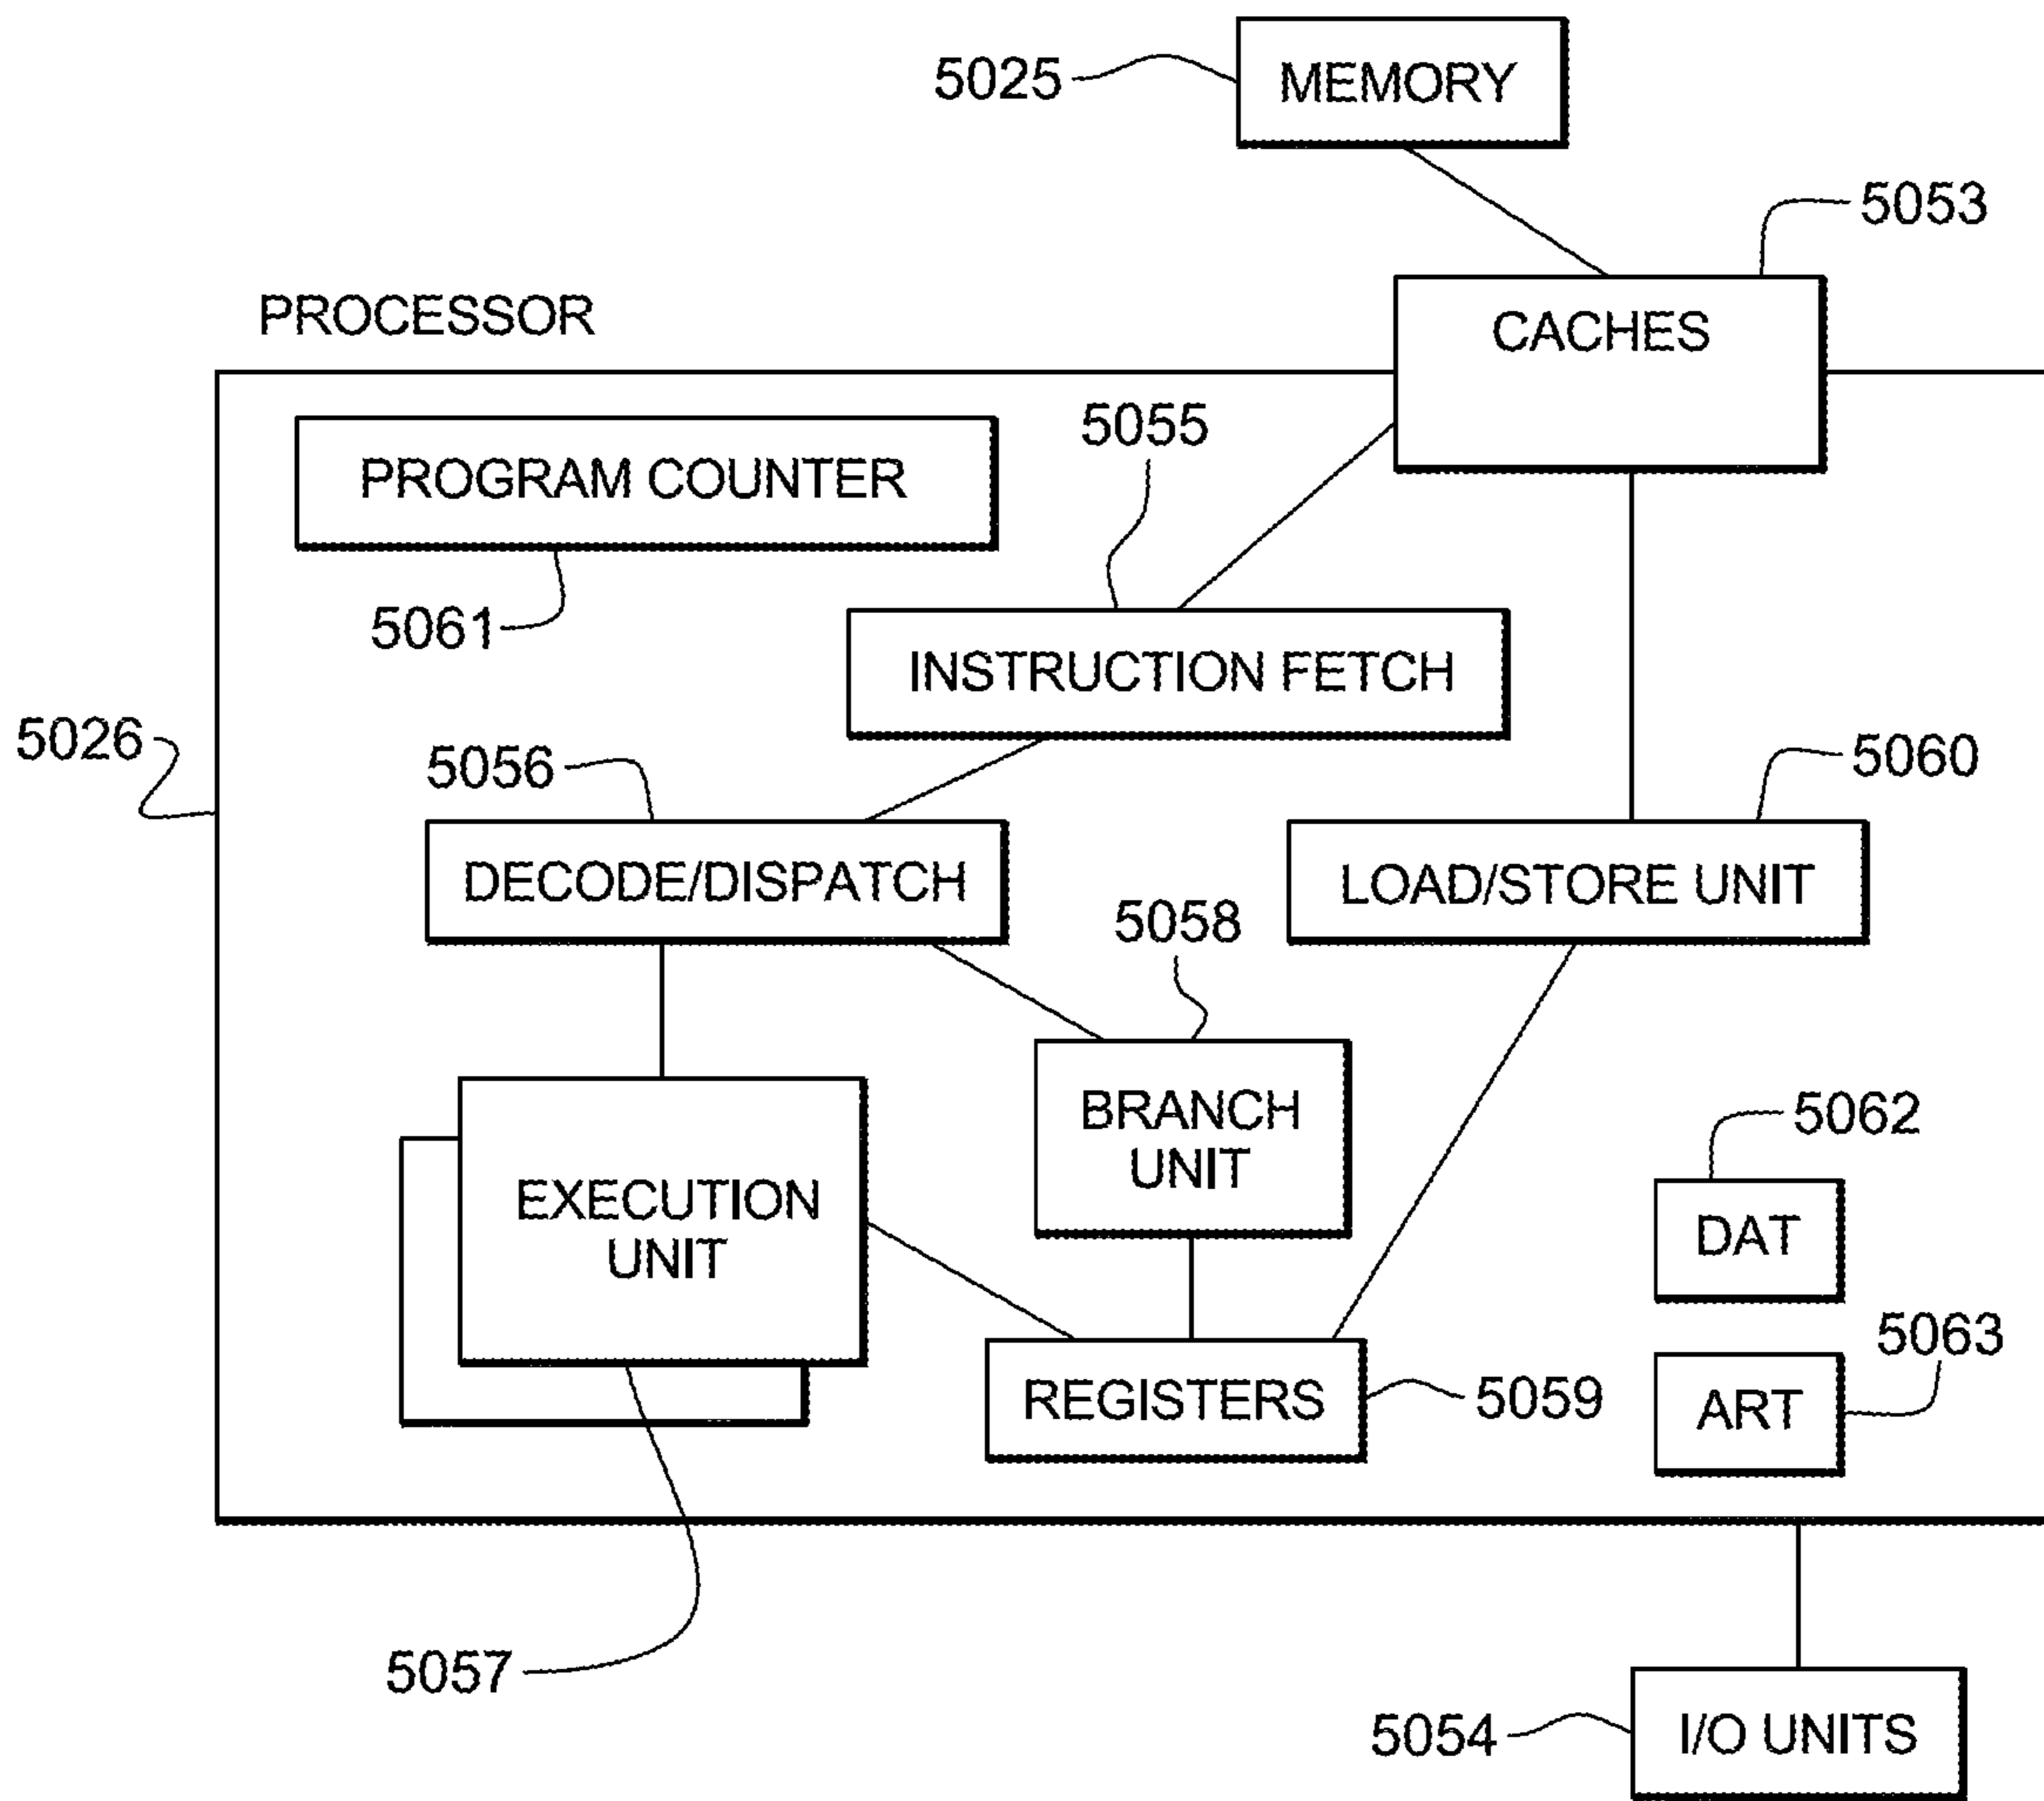


FIG. 19

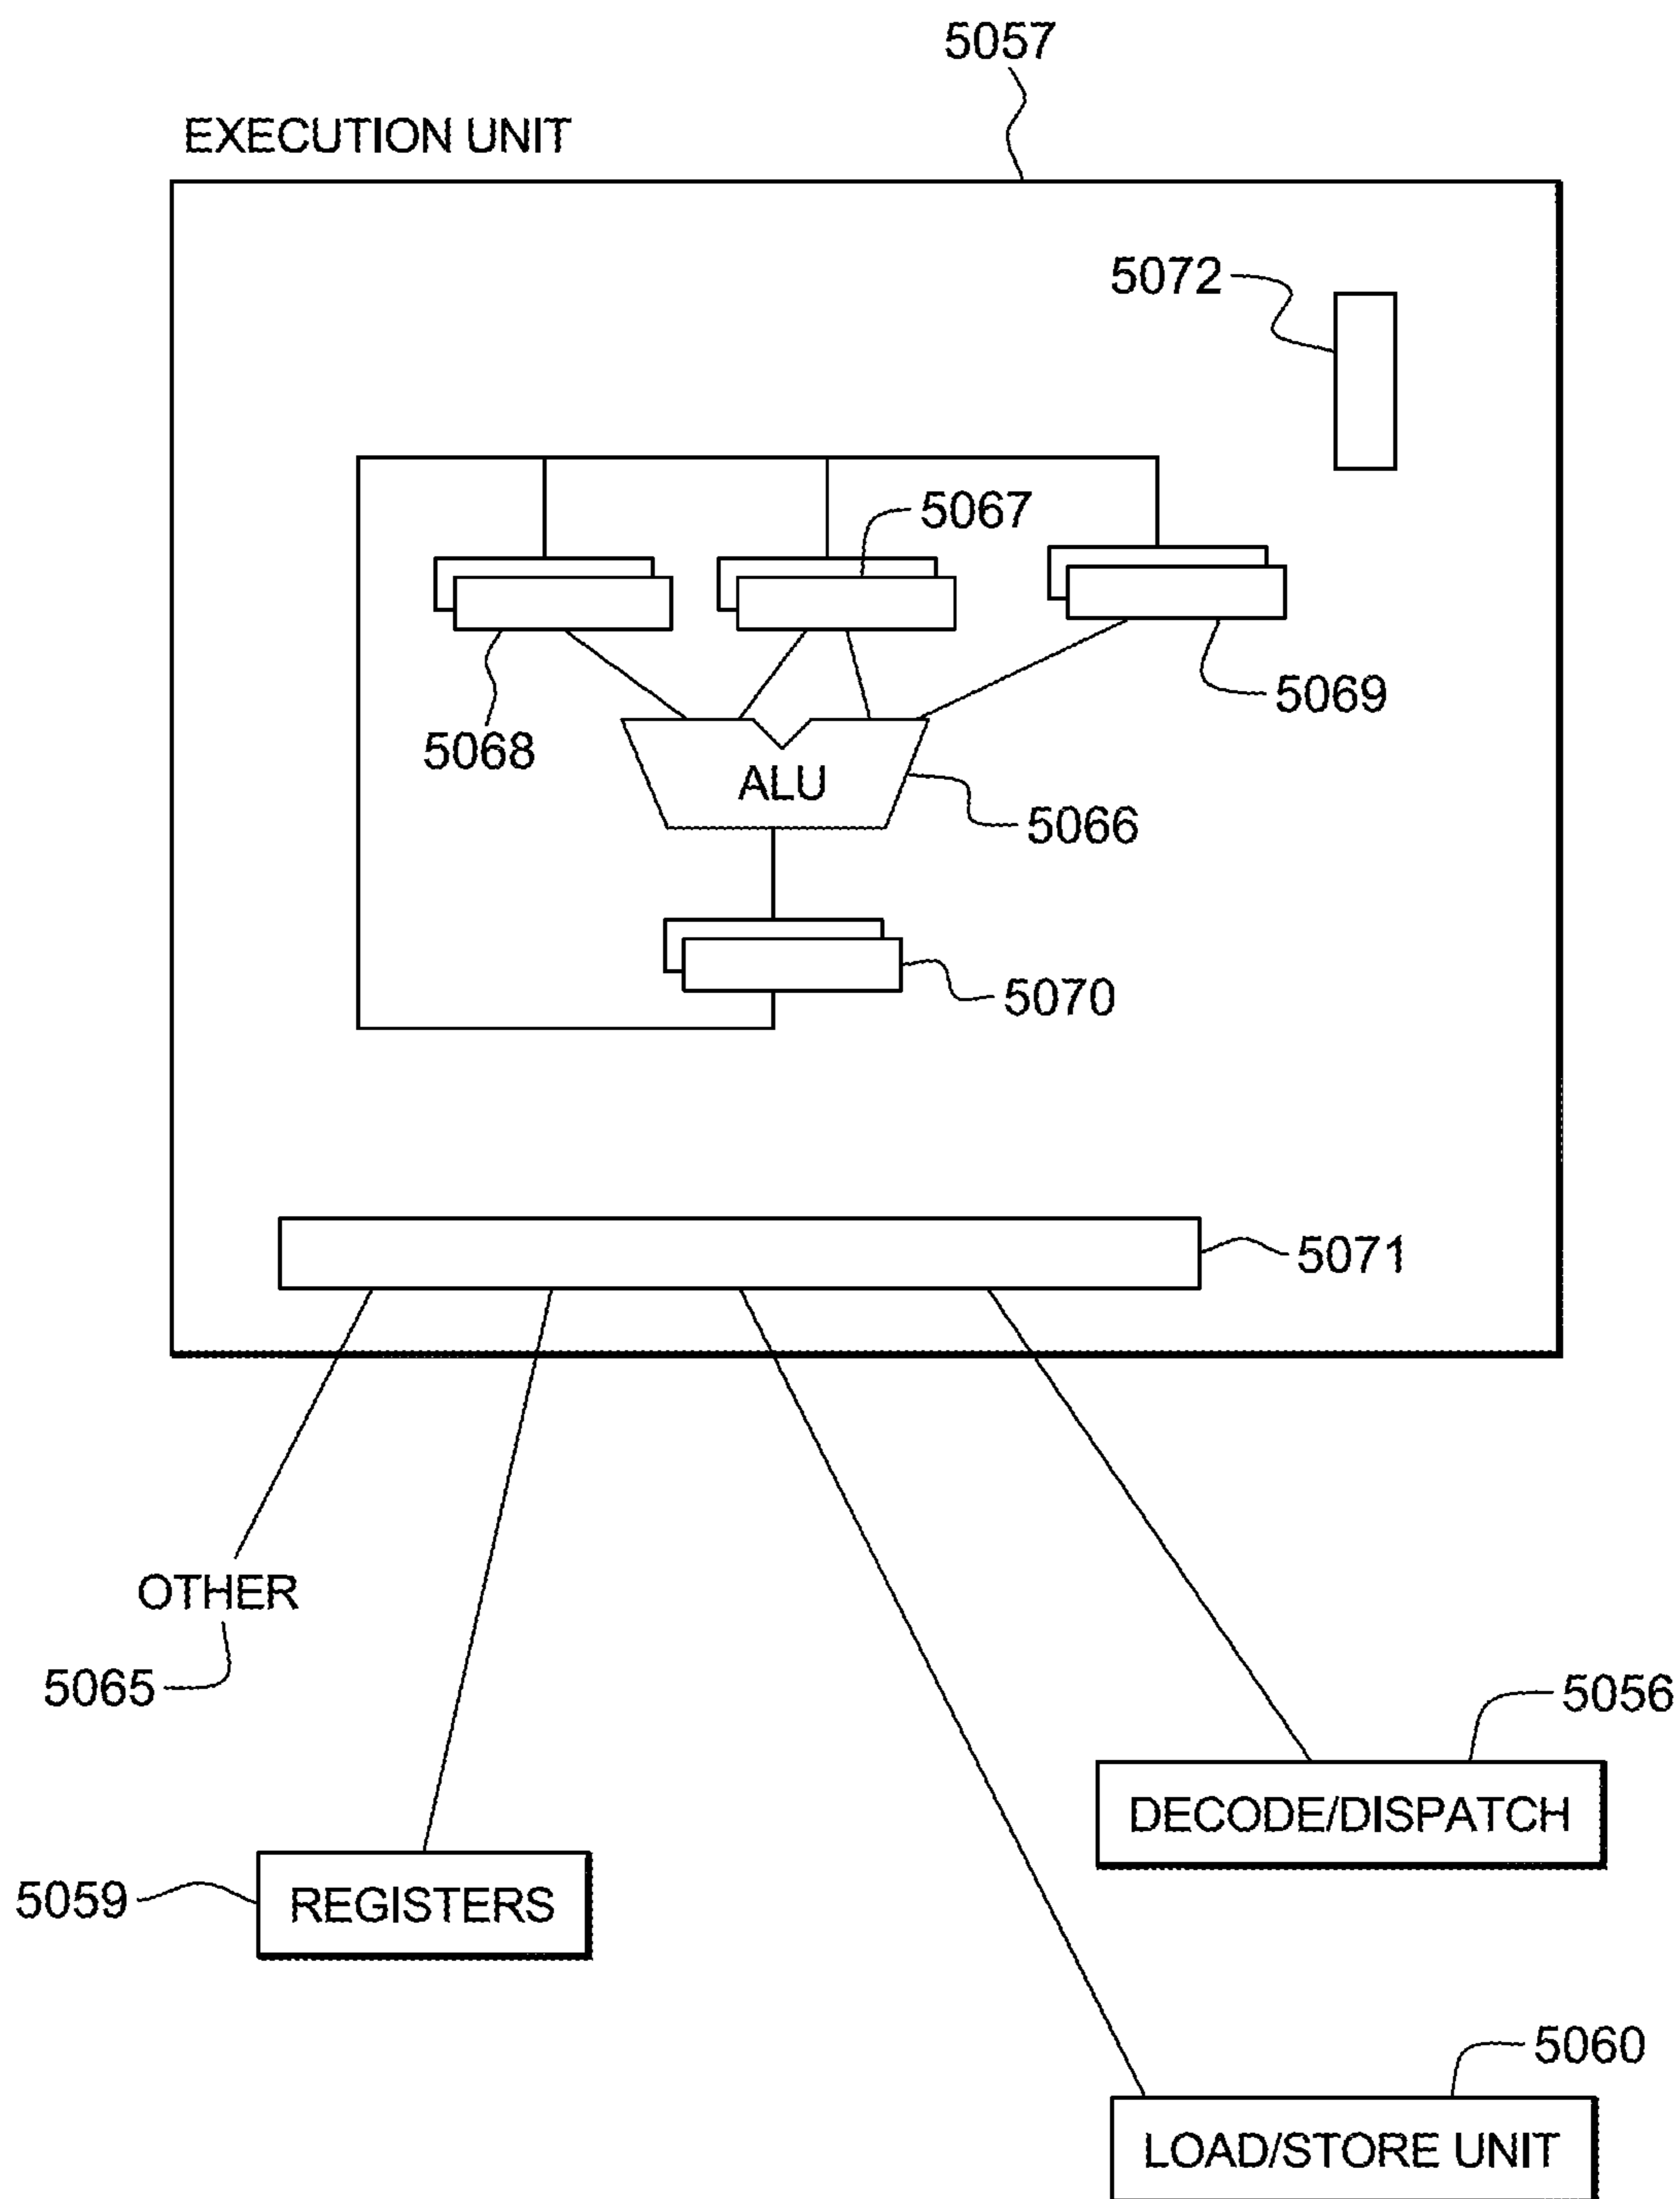


FIG. 20A

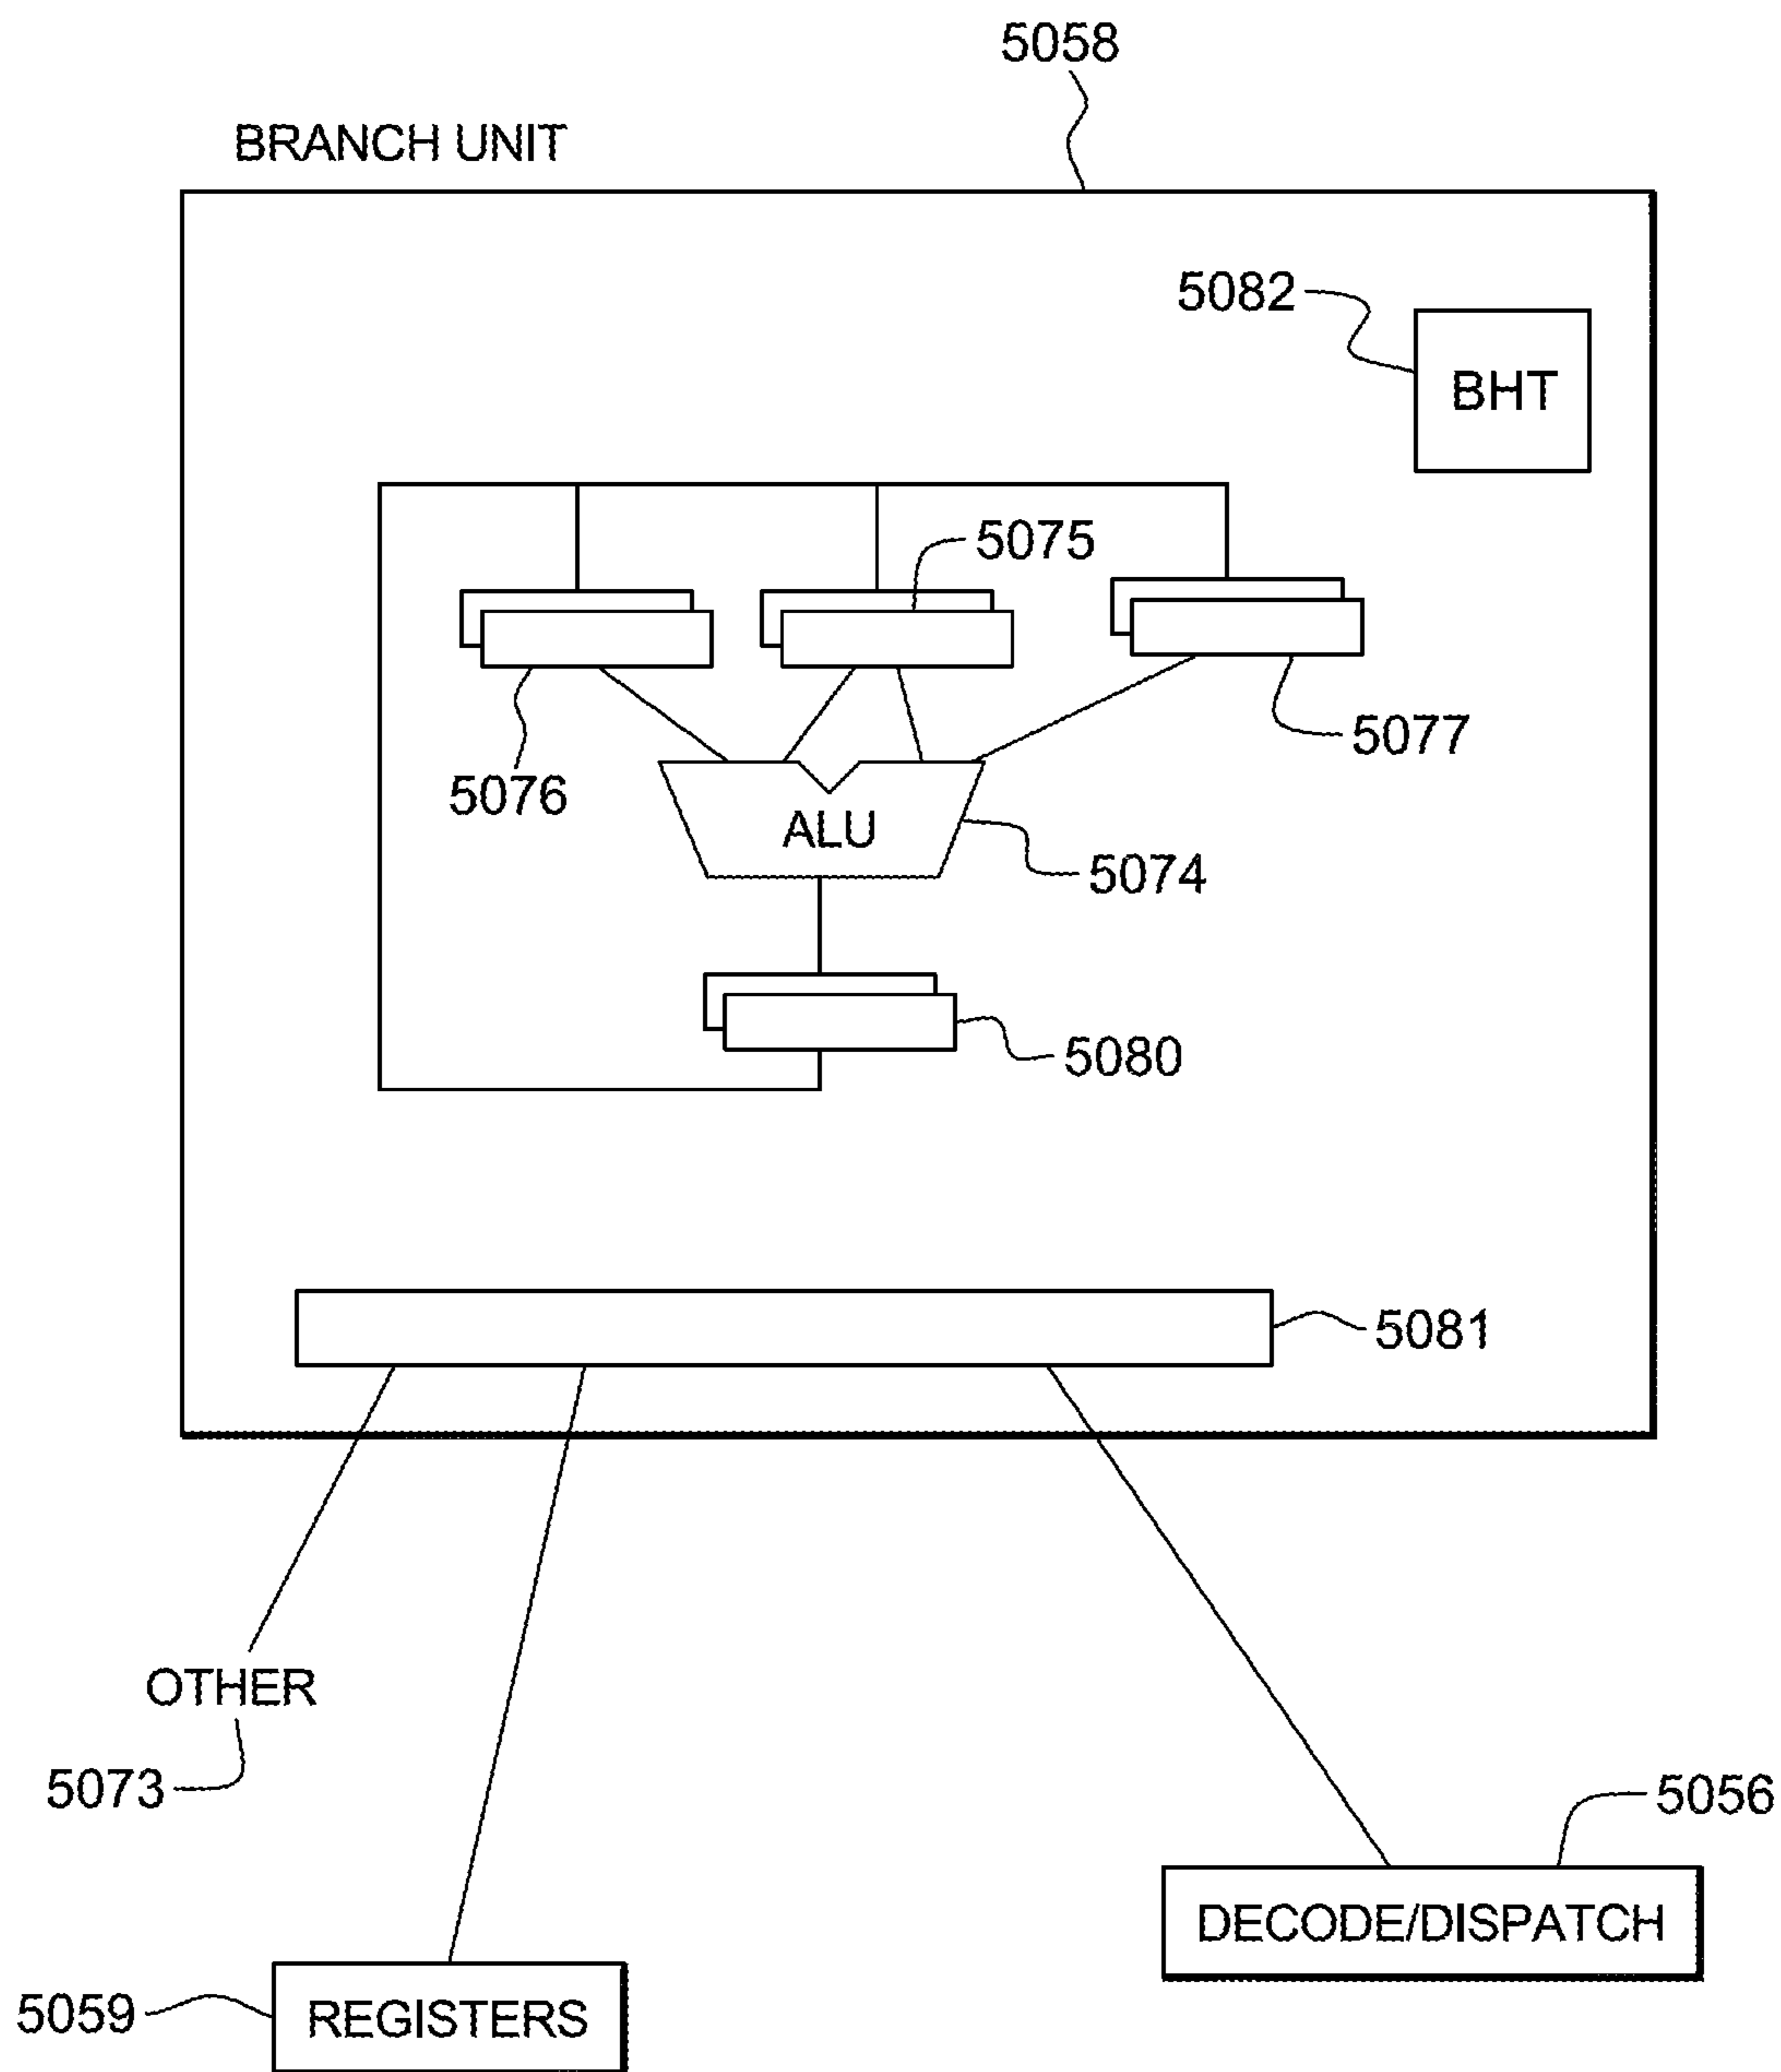


FIG. 20B

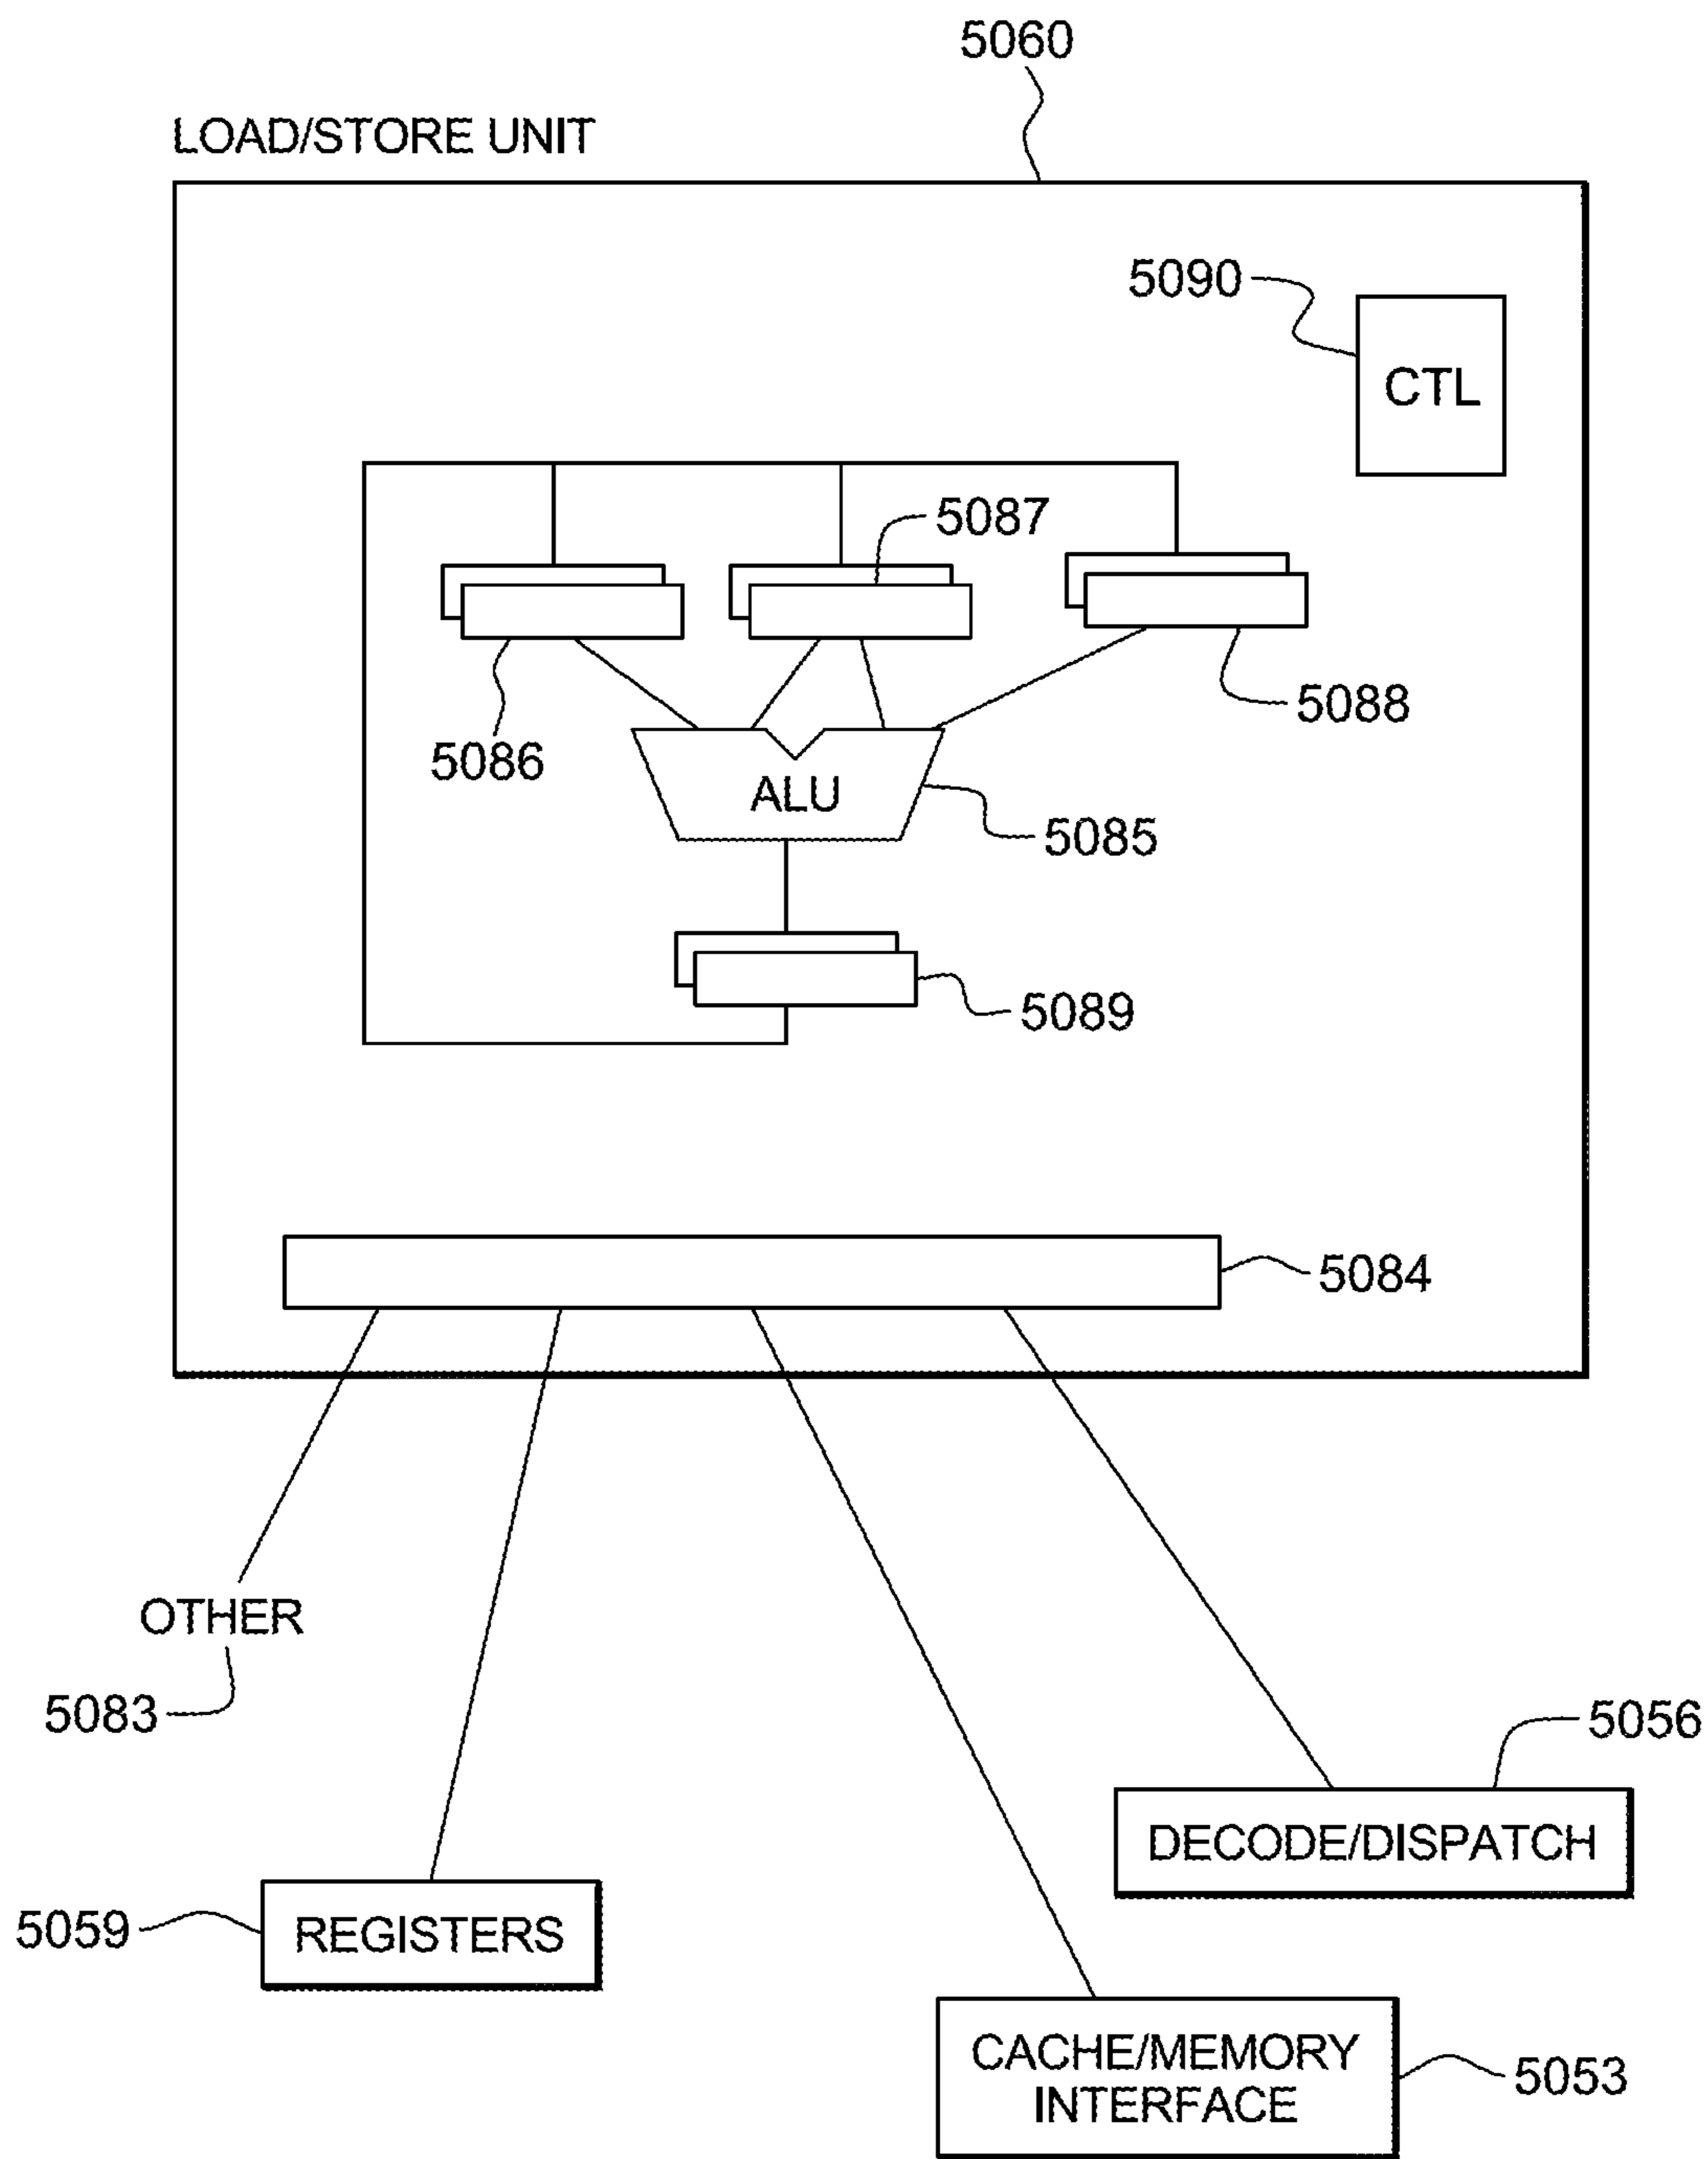


FIG. 20C

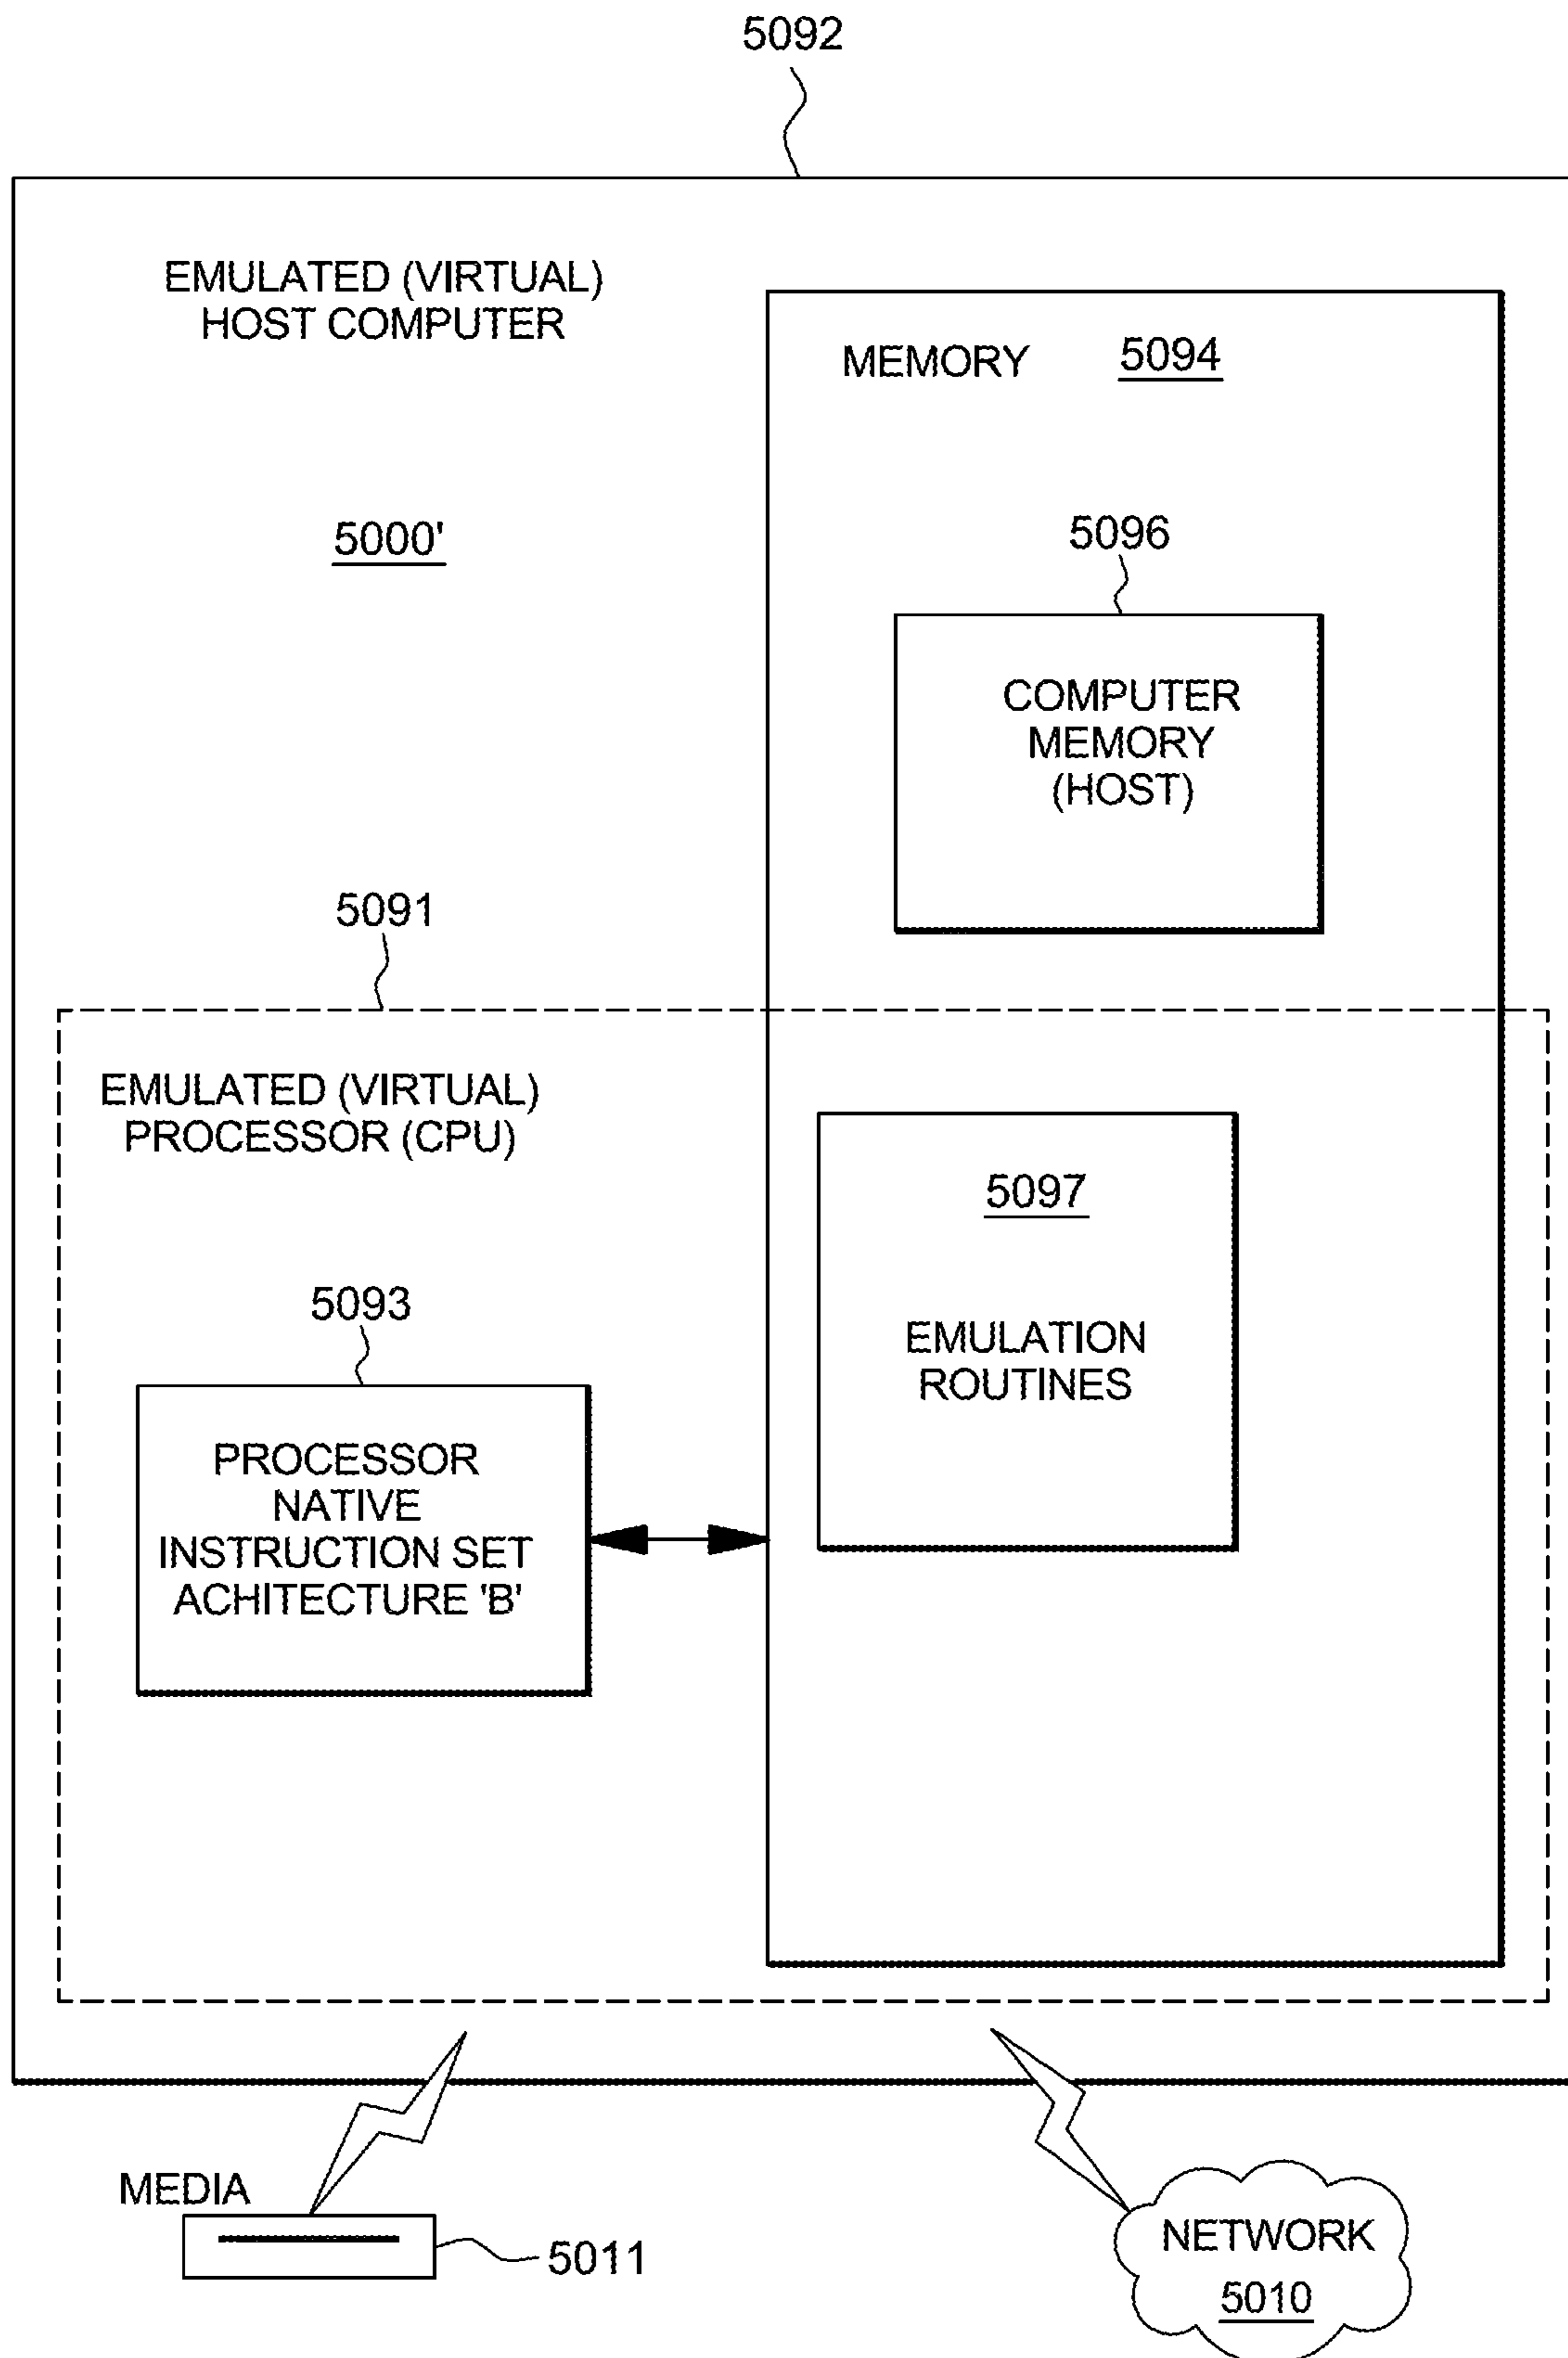


FIG. 21

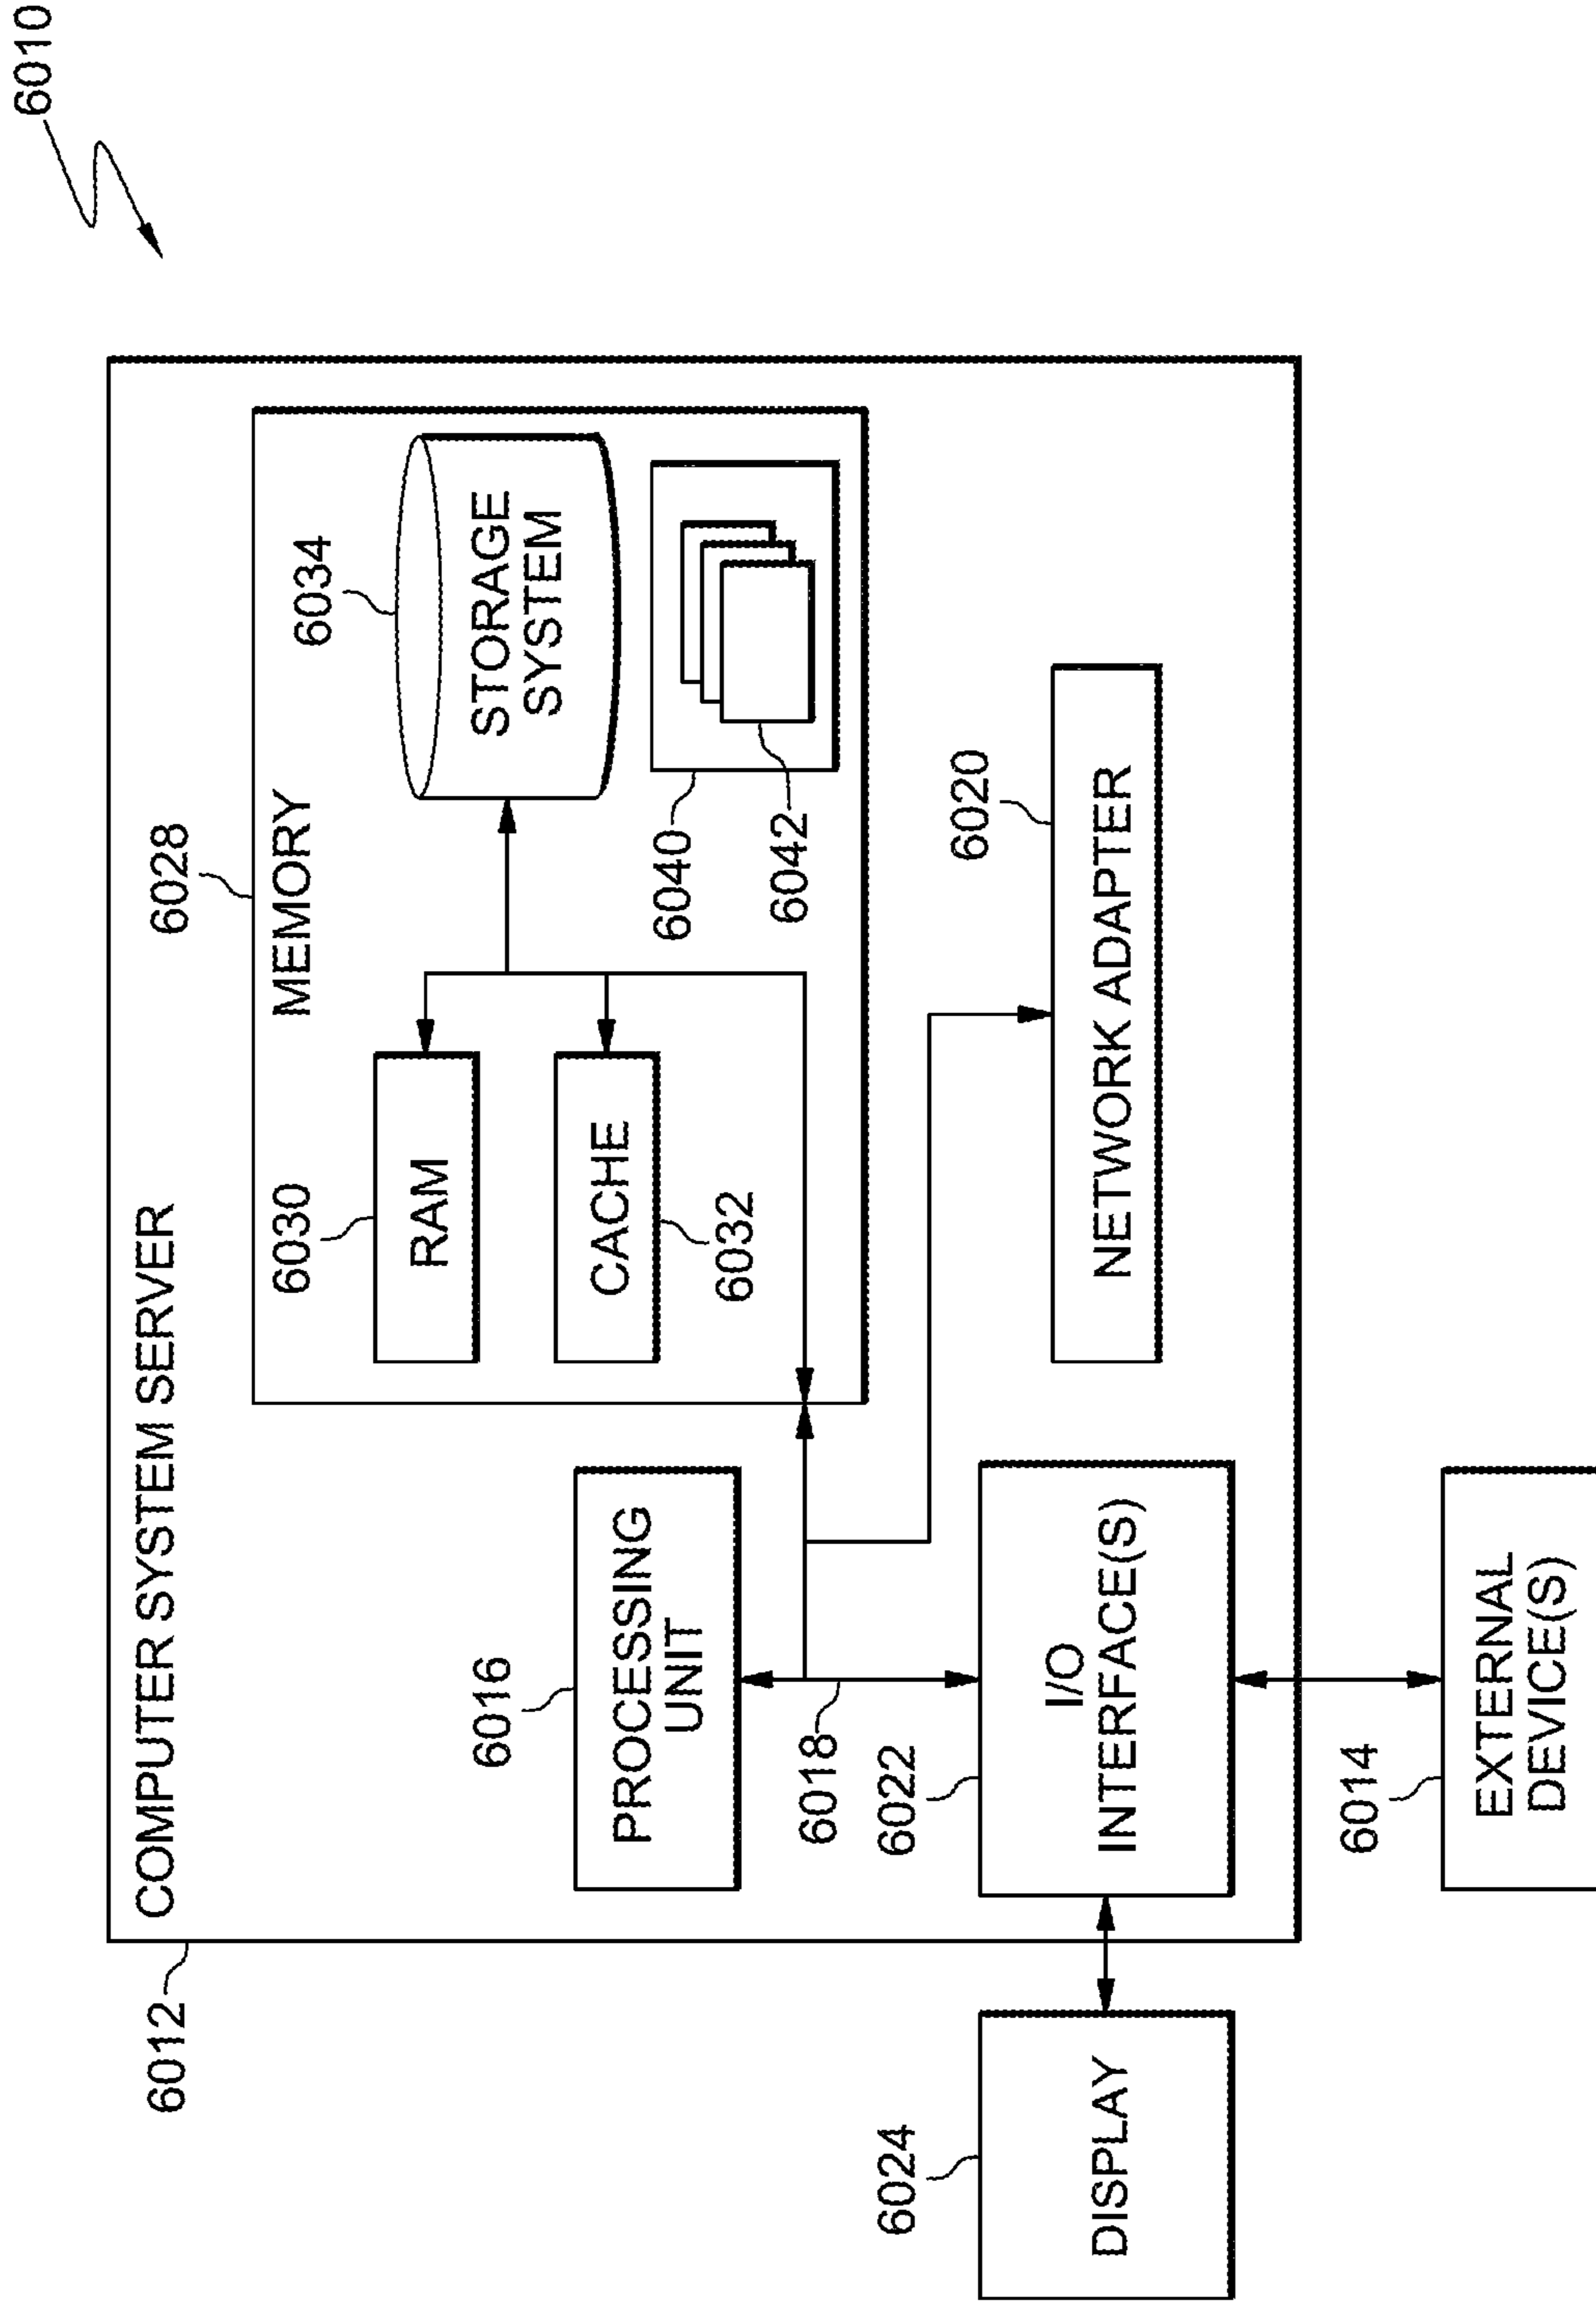


FIG. 22

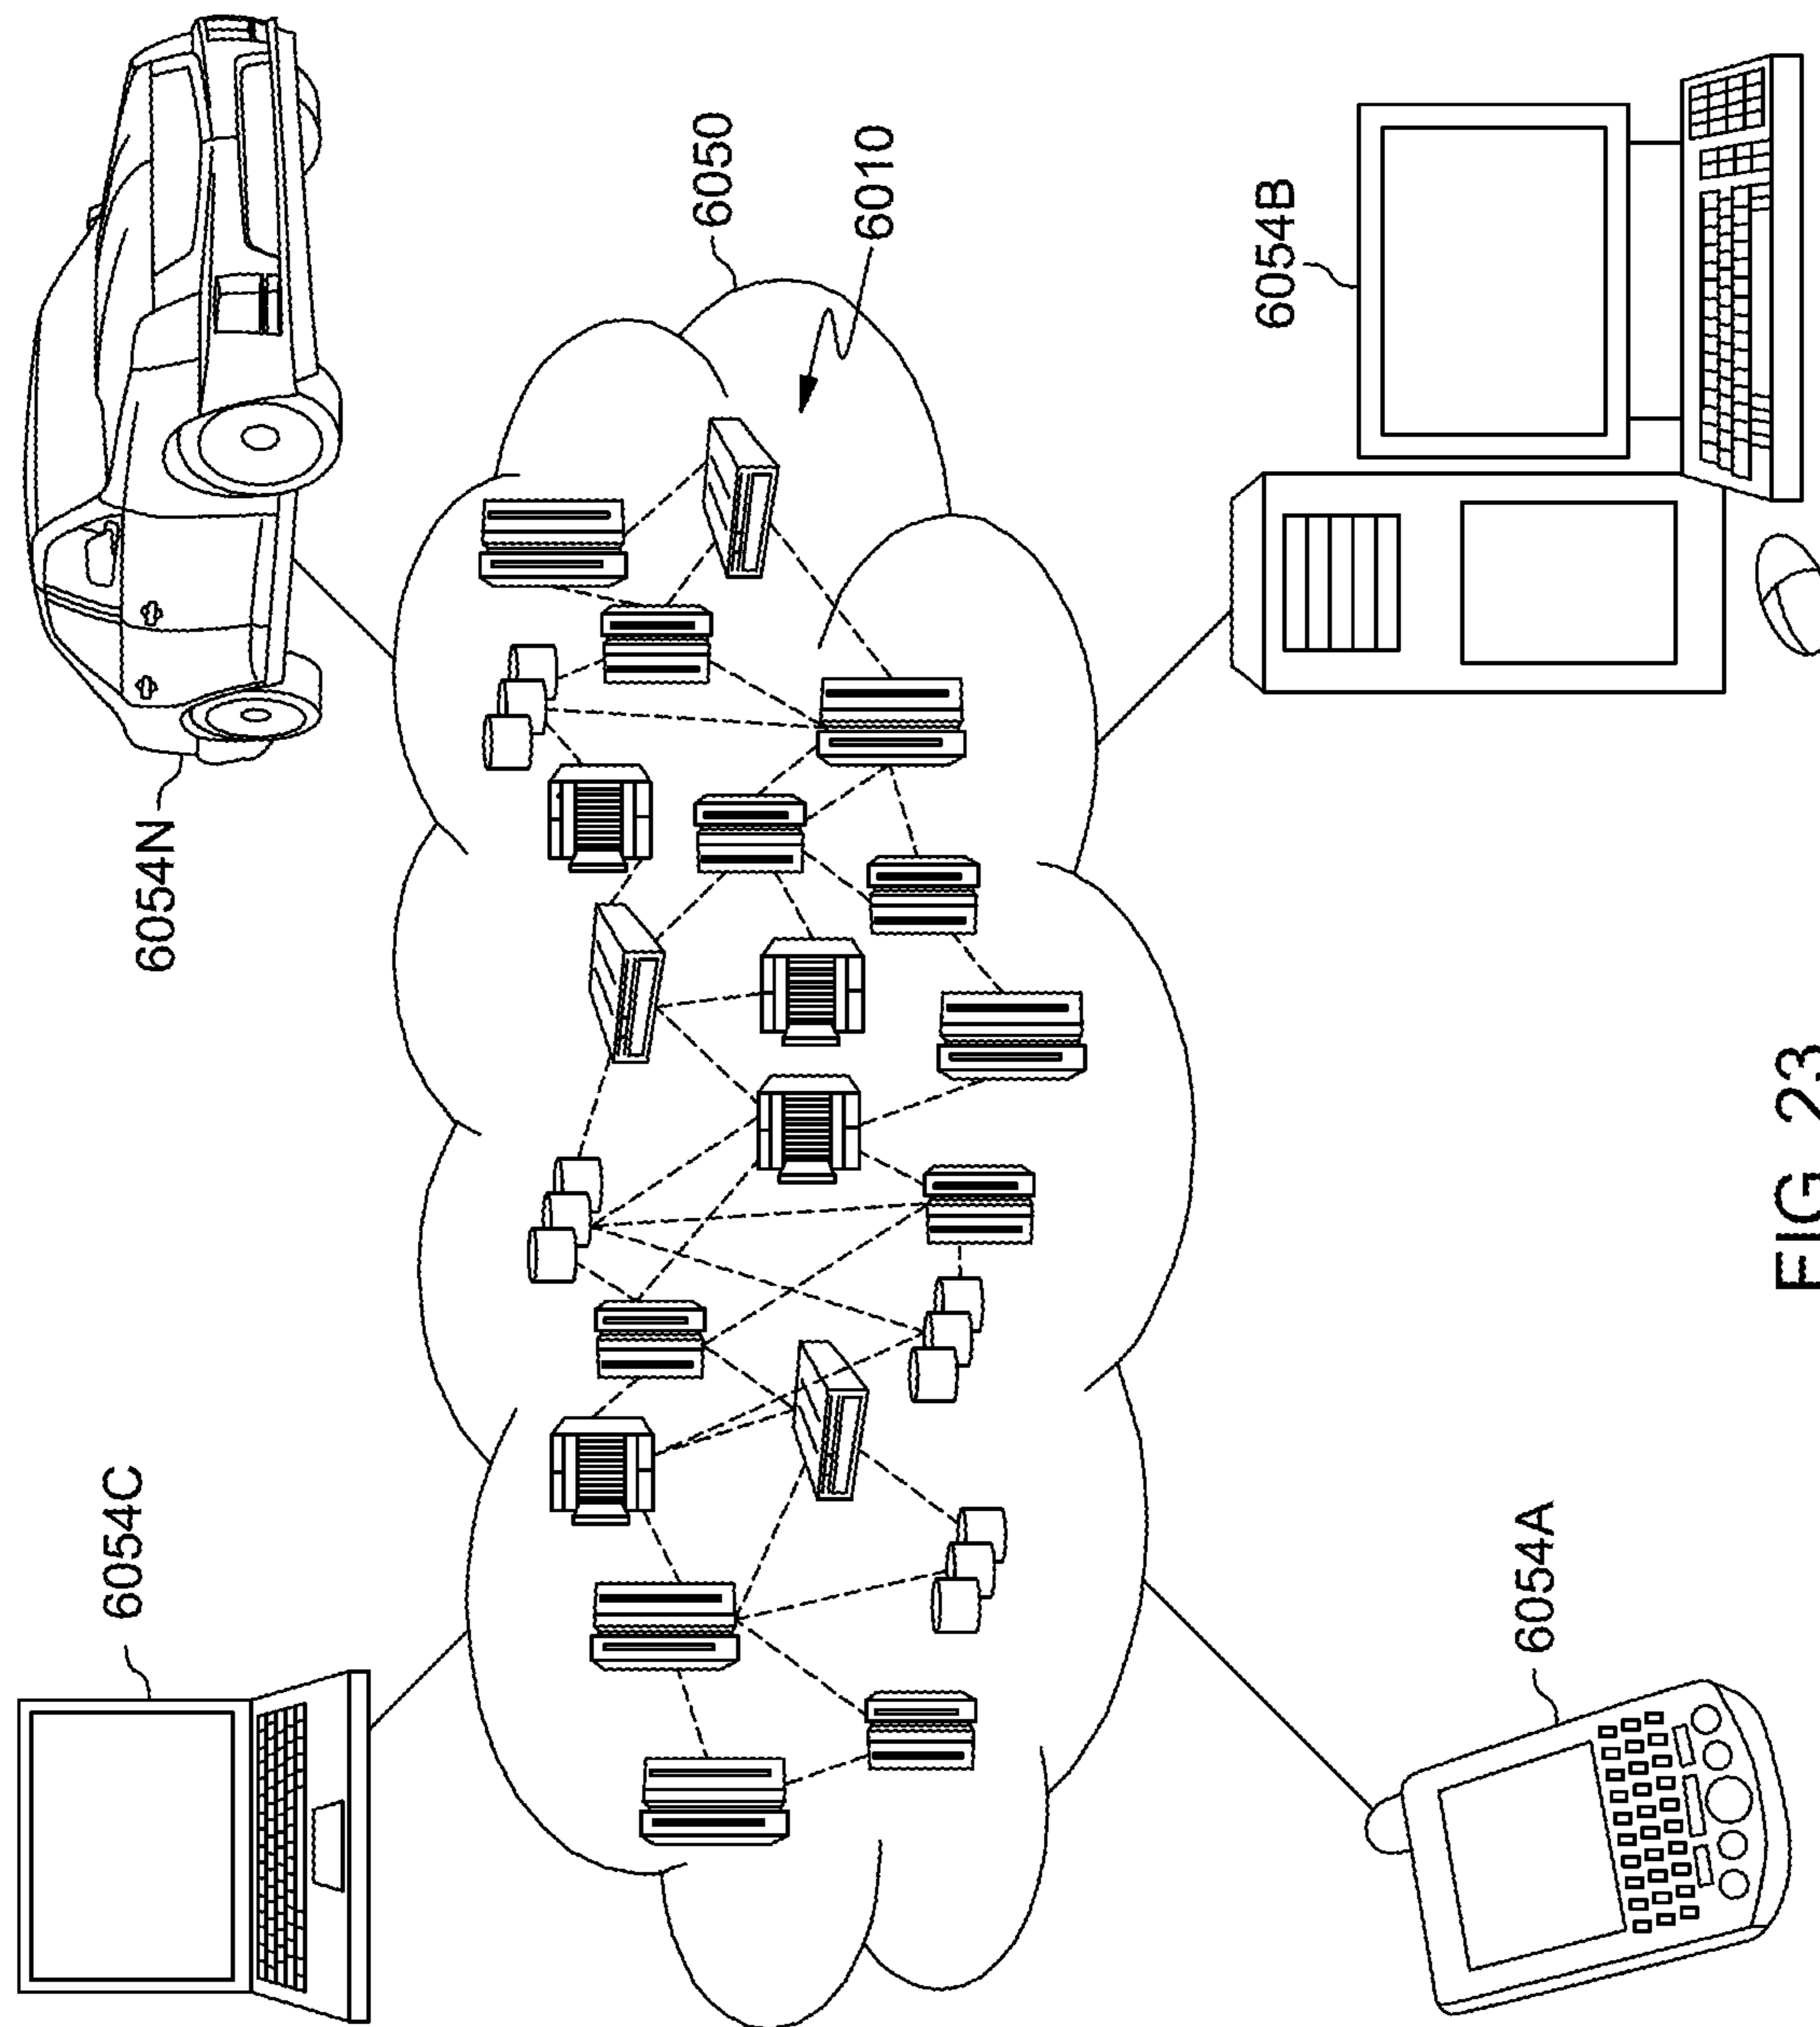


FIG. 23

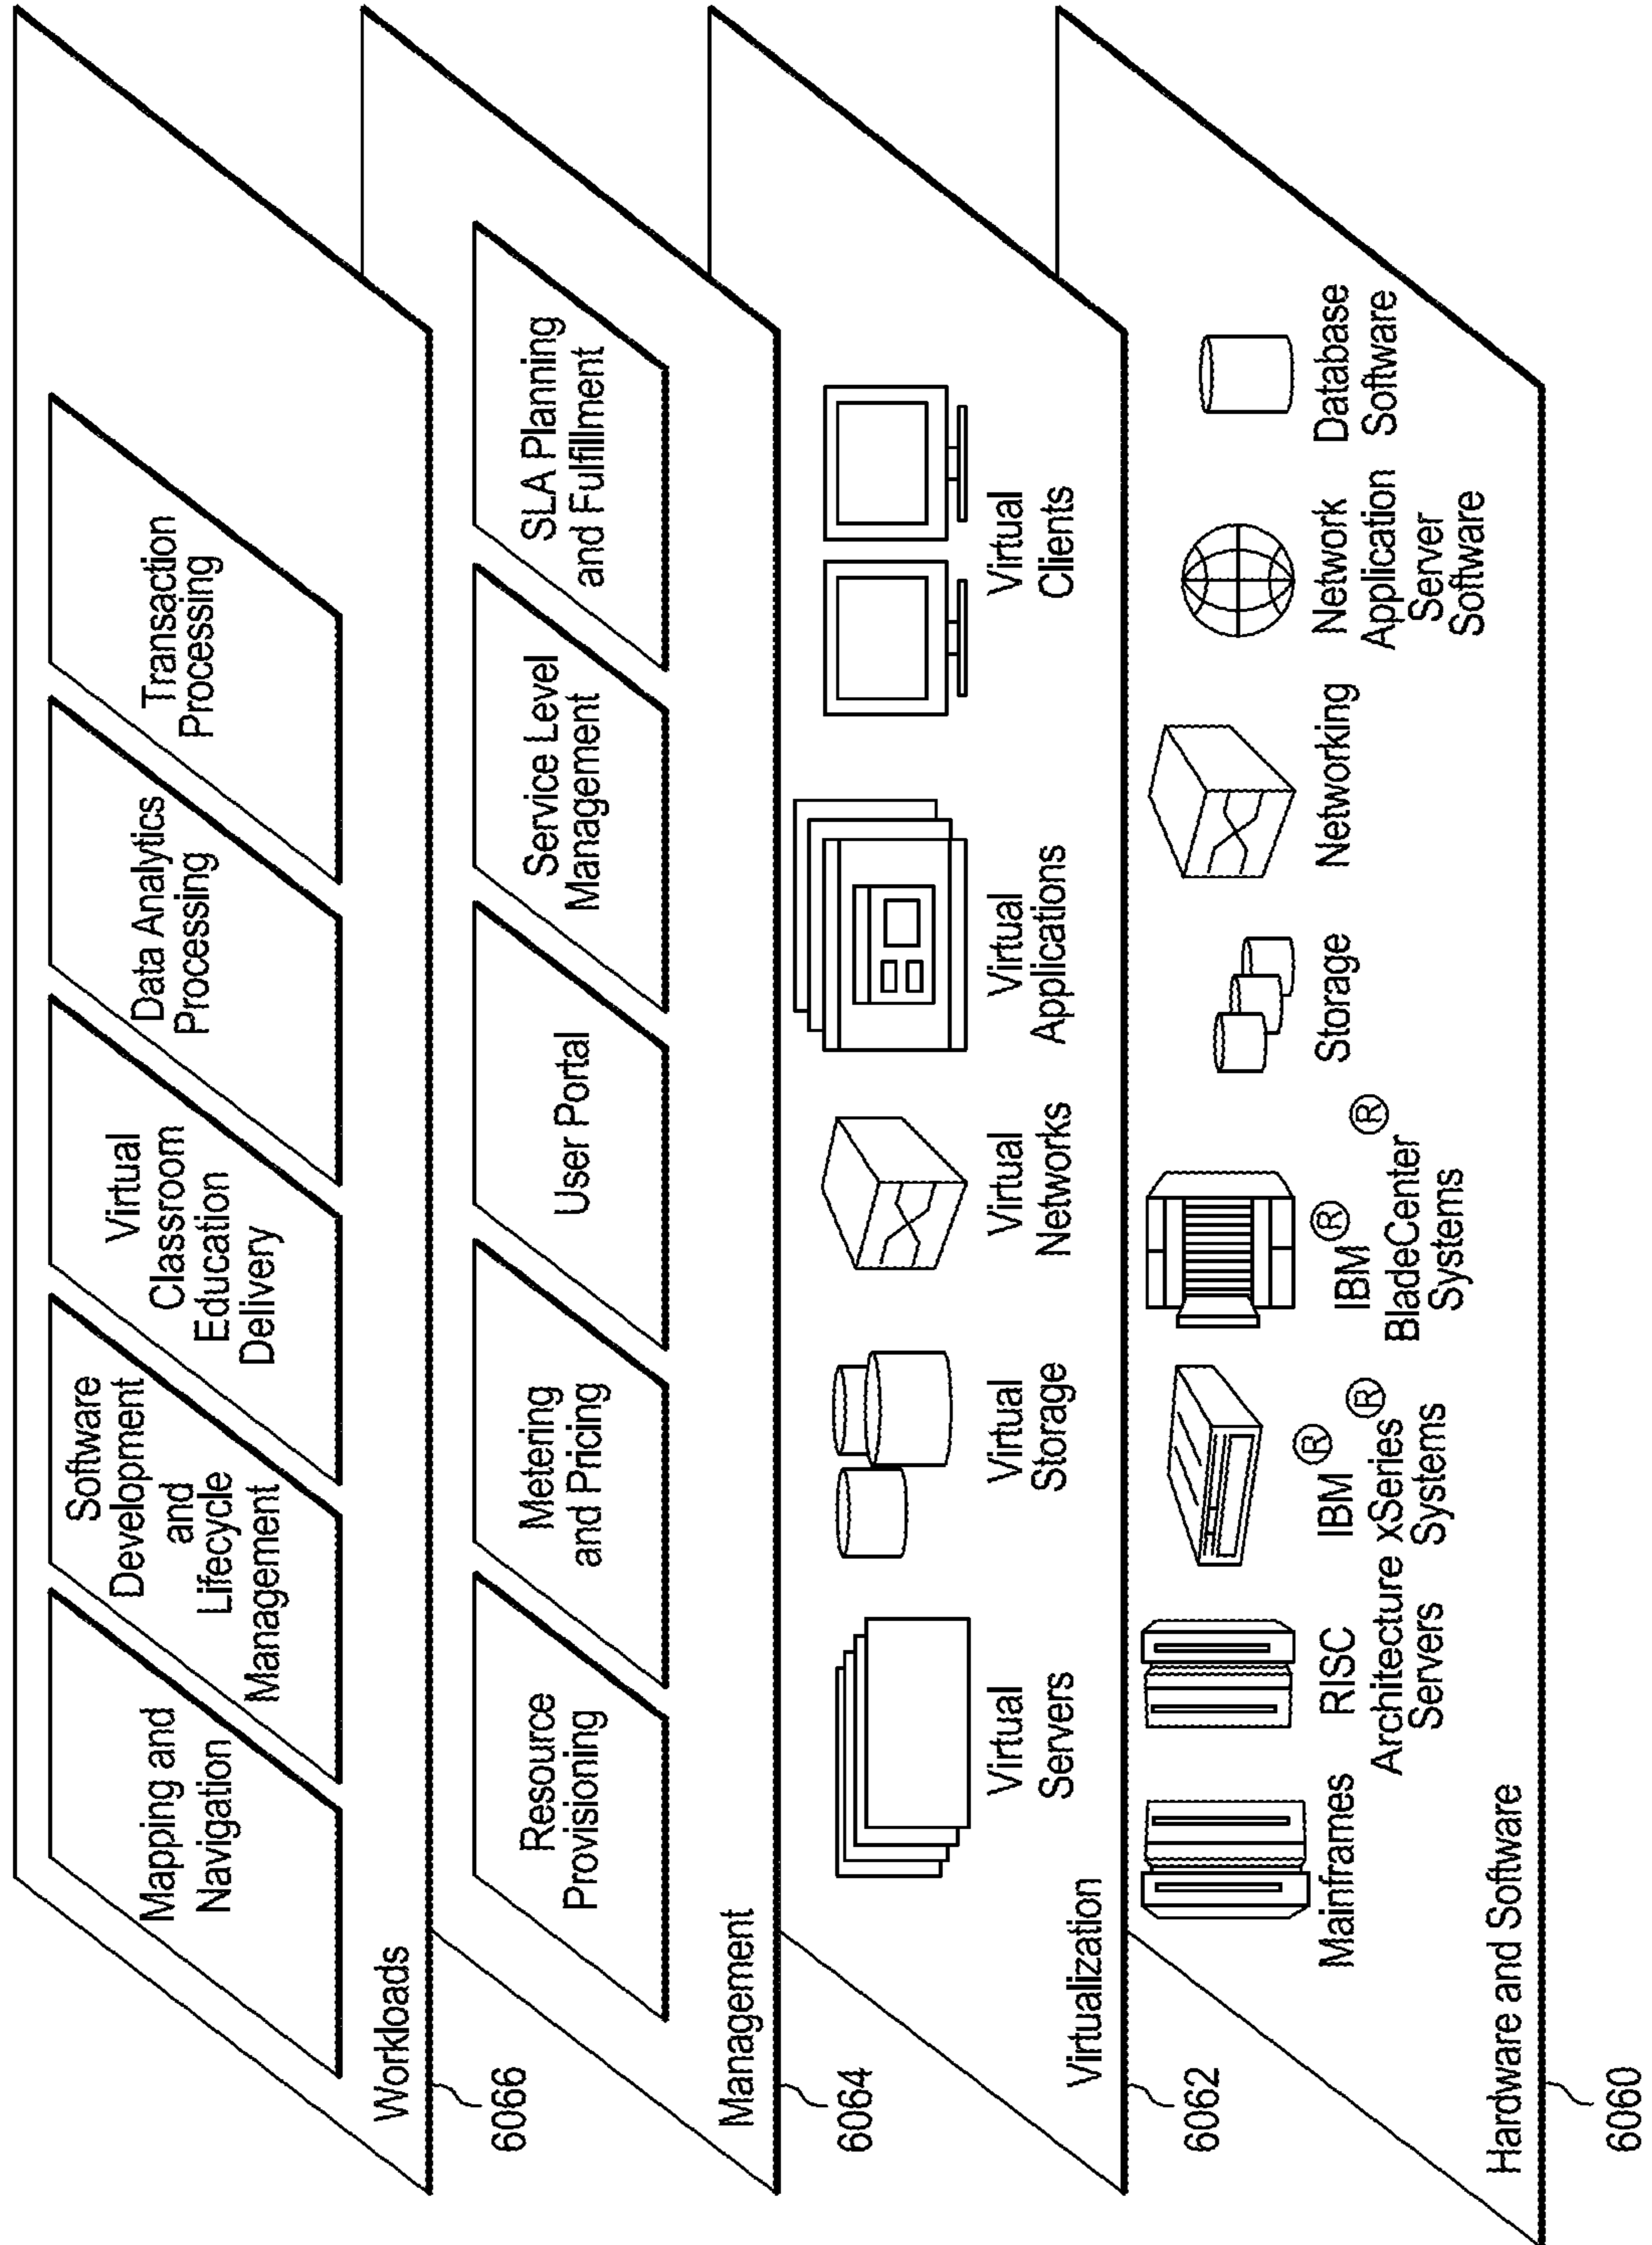


FIG. 24

INDICATOR-BASED PRIORITIZATION OF TRANSACTIONS

BACKGROUND

One or more aspects relate, in general, to multiprocessing computing environments, and in particular, to transactional processing within such computing environments.

In a transactional system, one or more processors perform transactions execution. During execution of a given transaction, a conflict may be detected and if a conflict is detected, the given transaction is aborted. Conflicts occur, for example, when one processor writes data that another processor is reading. A transactional core does not currently have any means to avoid or defer a conflict, thus, it aborts the transaction.

When constrained transactions are being executed, they complete in accordance with the Industry Standard Architecture (ISA) specification. The key difference between constrained and non-constrained transactions is that in the case of a constrained transaction, in the absence of repeated interruptions or other constraint violations, a constrained transaction is assured of eventual completion and therefore does not need a fall-back path. Thus, when a constrained transaction fails repeatedly, the computer system re-executes the transaction with escalating means to guarantee completion. The escalations employed include, but are not limited to, acquiring locks to exclude other cores. These escalations not only negatively impact the performance of the computer system as a whole, but they are provide no guarantee that the constrained transaction will be completed.

SUMMARY

Shortcomings of the prior art are overcome and additional advantages are provided through the provision of a computer program product for prioritizing transactions. The computer program product comprises a storage medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method. The method includes, for instance: initiating, by a processor, in a computing environment, an execution of a transaction, wherein the processor comprises a transactional core, and wherein the execution of the transaction is performed by the transactional core; obtaining, by the processor, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional core in the computing environment; determining, by the processor, if the transactional core comprises an indicator; based on determining that the transactional core comprises an indicator, ignoring, by the processor, the conflict and utilizing the transactional core to complete executing the transaction.

Shortcomings of the prior art are overcome and additional advantages are provided through the provision of a method of prioritizing transactions. The method includes, for instance: initiating, by a processor, in a computing environment, an execution of a transaction, wherein the processor comprises a transactional core, and wherein the execution of the transaction is performed by the transactional core; obtaining, by the processor, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional core in the computing environment; determining, by the processor, if the transactional core comprises an indicator; based on determining that the transactional core comprises an indi-

cator, ignoring, by the processor, the conflict and utilizing the transactional core to complete executing the transaction.

Shortcomings of the prior art are overcome and additional advantages are provided through the provision of a system for prioritizing transactions. The system includes a memory, one or more processor in communication with the memory, and program instructions executable by the one or more processor via the memory to perform a method. The method includes, for instance: initiating, by a processor, in a computing environment, an execution of a transaction, wherein the processor comprises a transactional core, and wherein the execution of the transaction is performed by the transactional core; obtaining, by the processor, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional core in the computing environment; determining, by the processor, if the transactional core comprises an indicator; based on determining that the transactional core comprises an indicator, ignoring, by the processor, the conflict and utilizing the transactional core to complete executing the transaction.

Methods and systems relating to one or more aspects are also described and claimed herein. Further, services relating to one or more aspects are also described and may be claimed herein.

Additional features and advantages are realized through the techniques described herein. Other embodiments and aspects are described in detail herein and are considered a part of the claimed aspects. For example, in an embodiment of the present invention, the transactions is a constrained transaction. In an embodiment of the present invention, the method also includes: monitoring, by the processor, the executing of the transaction by the transactional core; and based on the transactional core completing the executing by committing the transaction, clearing, by the processor, the indicator. The method in further embodiments of the present invention also includes: obtaining, by the processor, a request from the transactional core for the indicator; determining, by the processor, if a super core status is available; and based on determining that a super core status is available, modifying, by the processor, the transactional core, to provide the indicator.

In another embodiment of the present invention, the method also includes: based on determining that the transactional core does not comprise an indicator, aborting, by the processor, the transaction; and incrementing, by the processor, a value associated with the transaction. In this embodiment, the method may also include: obtaining, by the processor, the value associated with the transaction; evaluating, by the processor, the value to determine if the value meets a threshold; based on determining that the value meets the threshold, requesting, by the processor, an indicator for the transactional core; determining, by the processor, that a super core status is available; and based on determining, modifying, by the processor, the transactional core, to provide the indicator. In this embodiment, the determining if the super core status is available may include: accessing, by the processor, cores in the computing environment to determine if a threshold number of cores comprise indicators; and determining, by the processor, that the super core status is available based on less than the threshold number of cores comprising indicators.

In an embodiment of the present invention, the indication comprises a request from the additional core related to data utilized by the transactional core during the execution of the transaction. The data may be located in a cache of the transactional core and is marked as read data or write data.

3

In this embodiment of the present invention, wherein the ignoring may include: denying, by the processor, the additional core access to the data. The method may further include: queuing, by the processor, the request; and based on completing executing the transaction, fulfilling, by the processor, the request, by providing access to the data to the additional core.

BRIEF DESCRIPTION OF THE SEVERAL
VIEWS OF THE DRAWINGS

One or more aspects are particularly pointed out and distinctly claimed as examples in the claims at the conclusion of the specification. The foregoing and other objects, features, and advantages are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

FIG. 1 depicts one embodiment of a computing environment;

FIG. 2A depicts one example of a Transaction Begin (TBEGIN) instruction;

FIG. 2B depicts one embodiment of further details of a field of the TBEGIN instruction of FIG. 2A;

FIG. 3A depicts on example of a Transaction Begin Constrained (TBEGINC) instruction;

FIG. 3B depicts one embodiment of further details of a field of the TBEGINC instruction of FIG. 3A;

FIG. 4 depicts one example of a Transaction End (TEND) instruction;

FIG. 5 depicts one example of a Transaction Abort (TABORT) instruction;

FIG. 6 depicts one example of nested transactions;

FIG. 7 depicts one example of a transaction diagnostic block;

FIG. 8 depicts a workflow of a transaction as executing in a computing system.

FIG. 9 depicts an example of how program code may abort a transaction when there is a conflict, which is an issue addressed by aspects of some embodiments of the present invention;

FIG. 10 illustrates certain aspects of the functionality of a system where the super core indicator of certain embodiments of the present invention is utilized;

FIG. 11 illustrates aspects of certain embodiments of the present invention;

FIG. 12 illustrates aspects of certain embodiments of the present invention;

FIG. 13 provides an additional illustration of the aspects in FIG. 12;

FIG. 14 depicts a workflow illustrating certain aspects of some embodiments of the present invention;

FIG. 15 depicts an embodiment of a computer program product;

FIG. 16 depicts one embodiment of a host computer system;

FIG. 17 depicts a further example of a computer system;

FIG. 18 depicts another example of a computer system comprising a computer network;

FIG. 19 depicts one embodiment of various elements of a computer system;

FIG. 20A depicts one embodiment of the execution unit of the computer system of FIG. 19;

FIG. 20B depicts one embodiment of the branch unit of the computer system of FIG. 19;

FIG. 20C depicts one embodiment of the load/store unit of the computer system of FIG. 19;

4

FIG. 21 depicts one embodiment of an emulated host computer system;

FIG. 22 depicts one embodiment of a cloud computing node;

FIG. 23 depicts one embodiment of a cloud computing environment; and

FIG. 24 depicts one example of abstraction model layers.

DETAILED DESCRIPTION

The accompanying figures, in which like reference numerals refer to identical or functionally similar elements throughout the separate views and which are incorporated in and form a part of the specification, further illustrate the present invention and, together with the detailed description of the invention, serve to explain the principles of the present invention. As understood by one of skill in the art, the accompanying figures are provided for ease of understanding and illustrate aspects of certain embodiments of the present invention. The invention is not limited to the embodiments depicted in the figures.

As understood by one of skill in the art, program code, as referred to throughout this application, includes both software and hardware. For example, program code in certain embodiments of the present invention includes fixed function hardware, while other embodiments utilized a software-based implementation of the functionality described. Certain embodiments combine both types of program code.

Aspects of embodiments of the present invention provide for the prioritization of certain transactions by program code executing on one or more processor designating a one or more core in a computing environment as a super core (also referred to as a transactional super core). When a core is designated a super core, it will complete a constrained transaction because even if there is a conflict, for example, the super core may receive a request for data it contains in its cache which are marked as a read or write data, it will not abort a transaction. This designated super core will ignore the request for data, it will not provide access to the data, but will instead, continue to execute the transaction to its completion. In an embodiment of the present invention, a core may be marked as a super core if it, for example, fails a transaction one or more times. To assure completion of the transaction, program code may switch the core from a transactional core status into a super core for transaction status. In an embodiment of the present invention, when the transaction is complete, the core is restored to its transactional status from the super core status.

In an embodiment of the present invention, in order to extend super core status (or the possibility of this status) to processing cores within a computing environment, each core includes an indicator, for example, a bit, including but not limited to, a super user bit. In certain circumstances, program code can set the indicator to designate a core as a super core. As aforementioned, this program code can be both hardware and/or software. Thus, in some embodiments of the present invention, the program code that designates a super core is hardware logic. In others embodiments, the super core status is designated by at least one of microcode, millicode, and/or firmware code. In some embodiments of the present invention, an Operating System (OS) and/or a hypervisor designated a core a super core. In an embodiment of the present invention, a processor may execute an explicit instruction may set a core to super core status. In an embodiment of the present invention, program code utilizes an algorithm to select a core to become a super core.

In an embodiment of the present invention, program code designates a first transactional core to be a super core at the beginning of a given transaction. Having been designated a super core, the first transactional core will start executing the transaction in the same manner as core that had not been designated a super core, and will start issuing read and write memory commands. However, because it is a super core, the first transactional core will not respect coherence protocol requests from other cores. For example, if a second transactional core requests data which the first transactional core has saved in its cache, the first transactional core will not respond with the data and status change, as per coherence protocol, but will continue executing the transaction until the transaction is complete.

In an embodiment of the present invention, program code initiates the execution of a transaction by a transactional core. The program code obtains, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional core in the computing environment. The program code determines if the transactional core comprises an indicator (e.g., the aforementioned super core indicator) and based on determining that the transactional core comprises an indicator, the program code ignores the conflict and utilizes the transactional core to complete executing the transaction.

In some embodiments of the present invention, the indication includes a request from the additional core related to data utilized by the transactional core during the execution of the transaction. In an embodiment of the present invention, the data is located in a cache of the transactional core and is marked as read data or write data. In an embodiment of the present invention, the program code ignores the conflict by denying the additional core access to the data. The transaction may be, for example, a constrained transaction.

In an embodiment of the present invention, where the indication includes the aforementioned request, the program code also queues the request and based on completing executing the transaction, fulfills the request, by providing the additional core with access to the data.

In an embodiment of the present invention, the program code may also monitor the executing of the transaction by the transactional core and based on the transactional core completing the executing (i.e., by committing the transaction), the program code may clear the indicator.

In an embodiment of the present invention, the program code obtains a request from the transactional core for the indicator. The program code determines if a super core status is available, and based on determining that a super core status is available, the program code modifies, the transactional core, to provide the indicator.

In an embodiment of the present invention, based on determining that the transactional core does not comprise an indicator, the program code aborts the transaction and increments value associated with the transaction. The value can be utilized in some embodiments of the present invention. For example, in an embodiment of the present invention, the program code obtains the value associated with the transaction and evaluates the value to determine if the value meets a threshold. Based on determining that the value meets the threshold, the program code requests an indicator for the transactional core. The program code then determines that a super core status is available; and based on determining that the status is available, modifies the transactional core, to provide the indicator. In an embodiment of the present invention, the program code determined that the

super core status is available by accessing cores in the computing environment to determine if a threshold number of cores comprise indicators and determining that the super core status is available based on less than the threshold number of cores comprising indicators.

Prior to describing the super core designation in detail, however, details regarding the transactional execution facility, including nonconstrained and constrained transactional execution modes, are discussed. As noted above, certain aspects of embodiments the present invention are particularly relevant to constrained transactions because even if aborted, the one or more processor will continue to attempt to execute these instructions, resulting in inefficiencies throughout the system and by indicating that a certain core executing the transaction is a super core, the transaction will be completed, rather than repeatedly attempted, despite conflicts that would have impeded completion without the super code designation.

The transactional execution facility introduces a CPU state called the transactional execution (TX) mode. Following a CPU reset, the CPU is not in the TX mode. The CPU enters the TX mode by a TRANSACTION BEGIN instruction, and leaves the TX mode by either (a) an outermost TRANSACTION END instruction (more details on inner and outer to follow), (b) a CONDITIONAL TRANSACTION END instruction that sets the condition code to 0; or (c) the transaction being aborted. While in the TX mode, storage accesses by the CPU appear to be block-concurrent as observed by other CPUs and the I/O subsystem. The storage accesses are either (a) committed to storage when the outermost transaction ends without aborting (i.e., e.g., updates made in a cache or buffer local to the CPU are propagated and stored in real memory and visible to other CPUs), or (b) discarded if the transaction is aborted.

Transactions may be nested. That is, while the CPU is in the TX mode, it may execute another TRANSACTION BEGIN instruction. The instruction that causes the CPU to enter the TX mode is called the outermost TRANSACTION BEGIN; similarly, the program is said to be in the outermost transaction. Subsequent executions of TRANSACTION BEGIN are called inner instructions; and the program is executing an inner transaction. The model provides a minimum nesting depth and a model-dependent maximum nesting depth. An EXTRACT TRANSACTION NESTING DEPTH instruction returns the current nesting depth value, and in a further embodiment, may return a maximum nesting-depth value. This technique uses a model called "flattened nesting" in which an aborting condition at any nesting depth causes all levels of the transaction to be aborted, and control is returned to the instruction following the outermost TRANSACTION BEGIN.

During processing of a transaction, a transactional access made by one CPU is said to conflict with either (a) a transactional access or nontransactional access made by another CPU, or (b) a nontransactional access made by the I/O subsystem, if both accesses are to any location within the same cache line, and one or both of the accesses is a store. In other words, in order for transactional execution to be productive, the CPU is not to be observed making transactional accesses until it commits. This programming model may be highly effective in certain environments; for example, the updating of two points in a doubly-linked list of a million elements. However, it may be less effective, if there is a lot of contention for the storage locations that are being transactionally accessed.

In one model of transactional execution (referred to herein as a nonconstrained transaction), when a transaction is

aborted, the program may either attempt to re-drive the transaction in the hopes that the aborting condition is no longer present, or the program may “fall back” to an equivalent non-transactional path. In another model of transactional execution (referred to herein as a constrained transaction), an aborted transaction is automatically re-driven by the CPU; in the absence of constraint violations, the constrained transaction is assured of eventual completion.

When initiating a transaction, the program can specify various controls, such as (a) which general registers are restored to their original contents if the transaction is aborted, (b) whether the transaction is allowed to modify the floating-point-register context, including, for instance, floating point registers and the floating point control register, (c) whether the transaction is allowed to modify access registers (ARs), and (d) whether certain program-exception conditions are to be blocked from causing an interruption. If a nonconstrained transaction is aborted, various diagnostic information may be provided. For instance, the outermost TBEGIN instruction that initiates a nonconstrained transaction may designate a program specified transaction diagnostic block (TDB). Further, the TDB in the CPU’s prefix area or designated by the host’s state description may also be used if the transaction is aborted due to a program interruption or a condition that causes interpretative execution to end, respectively.

Indicated above are various types of registers. These are further explained in detail herein. General registers may be used as accumulators in general arithmetic and logical operations. In one embodiment, each register contains 64 bit positions, and there are 16 general registers. The general registers are identified by the numbers 0-15, and are designated by a four-bit R field in an instruction. Some instructions provide for addressing multiple general registers by having several R fields. For some instructions, the use of a specific general register is implied rather than explicitly designated by an R field of the instruction.

In addition to their use as accumulators in general arithmetic and logical operations, 15 of the 16 general registers are also used as base address and index registers in address generation. In these cases, the registers are designated by a four-bit B field or X field in an instruction. A value of zero in the B or X field specifies that no base or index is to be applied, and thus, general register 0 is not to be designated as containing a base address or index.

Floating point instructions use a set of floating point registers. The CPU has 16 floating point registers, in one embodiment. The floating point registers are identified by the numbers 0-15, and are designated by a four bit R field in floating point instructions. Each floating point register is 64 bits long and can contain either a short (32-bit) or a long (64-bit) floating point operand.

A floating point control (FPC) register is a 32-bit register that contains mask bits, flag bits, a data exception code, and rounding mode bits, and is used during processing of floating point operations.

Further, in one embodiment, the CPU has 16 control registers, each having 64 bit positions. The bit positions in the registers are assigned to particular facilities in the system, such as Program Event Recording (PER) (discussed below), and are used either to specify that an operation can take place or to furnish special information required by the facility. In one embodiment, for the transactional facility, CR0 (bits 8 and 9) and CR2 (bits 61-63) are used, as described below.

The CPU has, for instance, 16 access registers numbered 0-15. An access register consists of 32 bit positions con-

taining an indirect specification of an address space control element (ASCE). An address space control element is a parameter used by the dynamic address translation (DAT) mechanism to translate references to a corresponding address space. When the CPU is in a mode called the access register mode (controlled by bits in the program status word (PSW)), an instruction B field, used to specify a logical address for a storage operand reference, designates an access register, and the address space control element specified by the access register is used by DAT for the reference being made. For some instructions, an R field is used instead of a B field. Instructions are provided for loading and storing the contents of the access registers and for moving the contents of one access register to another.

Each of access registers 1-15 can designate any address space. Access register 0 designates the primary address space. When one of access registers 1-15 is used to designate an address space, the CPU determines which address space is designated by translating the contents of the access register. When access register 0 is used to designate an address space, the CPU treats the access register as designating the primary address space, and it does not examine the actual contents of the access register. Therefore, the 16 access registers can designate, at any one time, the primary address space and a maximum of 15 other spaces.

In one embodiment, there are multiple types of address spaces. An address space is a consecutive sequence of integer numbers (virtual addresses), together with the specific transformation parameters which allow each number to be associated with a byte location in storage. The sequence starts at zero and proceeds left to right.

In, for instance, the z/Architecture, when a virtual address is used by a CPU to access main storage (a.k.a., main memory), it is first converted, by means of dynamic address translation (DAT), to a real address, and then, by means of prefixing, to an absolute address. DAT may use from one to five levels of tables (page, segment, region third, region second, and region first) as transformation parameters. The designation (origin and length) of the highest-level table for a specific address space is called an address space control element, and it is found for use by DAT in a control register or as specified by an access register. Alternatively, the address space control element for an address space may be a real space designation, which indicates that DAT is to translate the virtual address simply by treating it as a real address and without using any tables.

DAT uses, at different times, the address space control elements in different control registers or specified by the access registers. The choice is determined by the translation mode specified in the current PSW. Four translation modes are available: primary space mode, secondary space mode, access register mode and home space mode. Different address spaces are addressable depending on the translation mode.

At any instant when the CPU is in the primary space mode or secondary space mode, the CPU can translate virtual addresses belonging to two address spaces—the primary address space and the second address space. At any instant when the CPU is in the access register mode, it can translate virtual addresses of up to 16 address spaces—the primary address space and up to 15 AR-specified address spaces. At any instant when the CPU is in the home space mode, it can translate virtual addresses of the home address space.

The primary address space is identified as such because it consists of primary virtual addresses, which are translated by means of the primary address space control element (ASCE). Similarly, the secondary address space consists of

secondary virtual addresses translated by means of the secondary ASCE; the AR specified address spaces consist of AR specified virtual addresses translated by means of AR specified ASCEs; and the home address space consists of home virtual addresses translated by means of the home ASCE. The primary and secondary ASCEs are in control registers 1 and 7, respectively. AR specified ASCEs are in ASN-second-table entries that are located through a process called access-register translation (ART) using control registers 2, 5 and 8. The home ASCE is in control register 13.

One embodiment of a computing environment to incorporate and use one or more aspects of the transactional facility described herein is described with reference to FIG. 1.

Referring to FIG. 1, in one example, a computing environment **100** is based on the z/Architecture, offered by International Business Machines (IBM®) Corporation, Armonk, N.Y. The z/Architecture is described in an IBM Publication entitled “z/Architecture—Principles of Operation,” Publication No. SA22-7932-09, 10th Edition, September 2012, which is hereby incorporated by reference herein in its entirety.

Z/ARCHITECTURE, IBM, and Z/OS and Z/VM (referenced below) are registered trademarks of International Business Machines Corporation, Armonk, N.Y. Other names used herein may be registered trademarks, trademarks or product names of International Business Machines Corporation or other companies.

As one example, computing environment **100** includes a central processor complex (CPC) **102** coupled to one or more input/output (I/O) devices **106** via one or more control units **108**. Central processor complex **102** includes, for instance, a processor memory **104** (a.k.a., main memory, main storage, central storage) coupled to one or more central processors (a.k.a., central processing units (CPUs)) **110**, and an input/output subsystem **111**, each of which is described below.

Processor memory **104** includes, for example, one or more partitions **112** (e.g., logical partitions), and processor firmware **113**, which includes a logical partition hypervisor **114** and other processor firmware **115**. One example of logical partition hypervisor **114** is the Processor Resource/System Manager (PRISM), offered by International Business Machines Corporation, Armonk, N.Y.

A logical partition functions as a separate system and has one or more applications **120**, and optionally, a resident operating system **122** therein, which may differ for each logical partition. In one embodiment, the operating system is the z/OS operating system, the z/VM operating system, the z/Linux operating system, or the TPF operating system, offered by International Business Machines Corporation, Armonk, N.Y. Logical partitions **112** are managed by logical partition hypervisor **114**, which is implemented by firmware running on processors **110**. As used herein, firmware includes, e.g., the microcode and/or millicode of the processor. It includes, for instance, the hardware-level instructions and/or data structures used in implementation of higher level machine code. In one embodiment, it includes, for instance, proprietary code that is typically delivered as microcode that includes trusted software or microcode specific to the underlying hardware and controls operating system access to the system hardware.

Central processors **110** are physical processor resources allocated to the logical partitions. In particular, each logical partition **112** has one or more logical processors, each of which represents all or a share of a physical processor **110** allocated to the partition. The logical processors of a par-

ticular partition **112** may be either dedicated to the partition, so that the underlying processor resource **110** is reserved for that partition; or shared with another partition, so that the underlying processor resource is potentially available to another partition. In one example, one or more of the CPUs include aspects of the transactional execution facility **130** and conditional transaction end facility **132** described herein.

Input/output subsystem **111** directs the flow of information between input/output devices **106** and main storage **104**. It is coupled to the central processing complex, in that it can be a part of the central processing complex or separate therefrom. The I/O subsystem relieves the central processors of the task of communicating directly with the input/output devices and permits data processing to proceed concurrently with input/output processing. To provide communications, the I/O subsystem employs I/O communications adapters. There are various types of communications adapters including, for instance, channels, I/O adapters, PCI cards, Ethernet cards, Small Computer Storage Interface (SCSI) cards, etc. In the particular example described herein, the I/O communications adapters are channels, and therefore, the I/O subsystem is referred to herein as a channel subsystem. However, this is only one example. Other types of I/O subsystems can be used.

The I/O subsystem uses one or more input/output paths as communication links in managing the flow of information to or from input/output devices **106**. In this particular example, these paths are called channel paths, since the communication adapters are channels.

The computing environment described above is only one example of a computing environment that can be used. Other environments, including but not limited to, non-partitioned environments, other partitioned environments, and/or emulated environments, may be used; embodiments are not limited to any one environment.

In accordance with one or more aspects, the transactional execution facility is an enhancement of the central processing unit that provides the means by which the CPU can execute a sequence of instructions—known as a transaction—that may access multiple storage locations, including the updating of those locations. As observed by other CPUs and the I/O subsystem, the transaction is either (a) completed in its entirety as a single atomic operation, or (b) aborted, potentially leaving no evidence that it ever executed (except for certain conditions described herein). Thus, a successfully completed transaction can update numerous storage locations without any special locking that is needed in the classic multiprocessing model.

The transactional execution facility includes, for instance, one or more controls; one or more instructions; transactional processing, including constrained and nonconstrained execution; and abort processing, each of which is further described below.

In one embodiment, three special purpose controls, including a transaction abort Program Status Word (PSW), a transaction diagnostic block (TDB) address, and a transaction nesting depth; five control register bits; and a plurality of general instructions, including TRANSACTION BEGIN (constrained and nonconstrained), TRANSACTION END, EXTRACT TRANSACTION NESTING DEPTH, TRANSACTION ABORT, and NONTRANSACTIONAL STORE are used to control the transactional execution facility. When the facility is installed, it is installed, for instance, in all CPUs in the configuration. A facility indication, bit 73 in one implementation, when one, indicates that the transactional execution facility is installed.

Further, in one aspect, when the transactional execution facility is installed, another facility, referred to as the conditional transaction end facility, may also be installed. The conditional transaction end facility is installed when, for instance, bit 55 of the facility indication bits is set to one. In one implementation, this bit is meaningful only when bit 73 representing the transactional execution facility is also one. When both the facilities are installed, then the CONDITIONAL TRANSACTION END instruction is also an enhancement of the CPU and is used to control transactional execution.

When the transactional execution facility is installed, the configuration provides a nonconstrained transactional execution facility, and optionally, a constrained transactional execution facility, each of which is described below. When facility indications 50 and 73, as examples, are both one, the constrained transactional execution facility is installed. Both facility indications are stored in memory at specified locations.

As used herein, the instruction name TRANSACTION BEGIN refers to the instructions having the mnemonics TBEGIN (Transaction Begin for a nonconstrained transaction) and TBEGINC (Transaction Begin for a constrained transaction). Discussions pertaining to a specific instruction are indicated by the instruction name followed by the mnemonic in parentheses or brackets, or simply by the mnemonic.

One embodiment of a format of a TRANSACTION BEGIN (TBEGIN) instruction is depicted in FIGS. 2A-2B. As one example, a TBEGIN instruction **200** (FIG. 2A) includes an opcode field **202** that includes an opcode specifying a transaction begin nonconstrained operation; a base field (B_1) **204**; a displacement field (D_1) **206**; and an immediate field (I_2) **208**. When the B_1 field is nonzero, the contents of the general register specified by B_1 **204** are added to D_1 **206** to obtain the first operand address.

When the B_1 field is nonzero, the following applies:

When the transaction nesting depth is initially zero, the first operand address designates the location of the 256 byte transaction diagnostic block, called the TBEGIN-specified TDB (described further below) into which various diagnostic information may be stored if the transaction is aborted. When the CPU is in the primary space mode or access register mode, the first operand address designates a location in the primary address space. When the CPU is in the secondary space or home space mode, the first operand address designates a location in the secondary or home address space, respectively. When DAT is off, the transaction diagnostic block (TDB) address (TDBA) designates a location in real storage.

Store accessibility to the first operand is determined. If accessible, the logical address of the operand is placed into the transaction diagnostic block address (TDBA), and the TDBA is valid.

When the CPU is already in the nonconstrained transactional execution mode, the TDBA is not modified, and it is unpredictable whether the first operand is tested for accessibility.

When the B_1 field is zero, no access exceptions are detected for the first operand and, for the outermost TBEGIN instruction, the TDBA is invalid.

The bits of the I_2 field are defined as follows, in one example:

General Register Save Mask (GRSM) **210** (FIG. 2B): Bits 0-7 of the I_2 field contain the general register save mask (GRSM). Each bit of the GRSM represents an even-odd pair

of general registers, where bit 0 represents registers 0 and 1, bit 1 represents registers 2 and 3, and so forth. When a bit in the GRSM of the outermost TBEGIN instruction is zero, the corresponding register pair is not saved. When a bit in the GRSM of the outermost TBEGIN instruction is one, the corresponding register pair is saved in a model dependent location that is not directly accessible by the program.

If the transaction aborts, saved register pairs are restored to their contents when the outermost TBEGIN instruction was executed. The contents of all other (unsaved) general registers are not restored when a transaction aborts.

The general register save mask is ignored on all TBEGINs except for the outermost one.

Allow AR Modification (A) **212**: The A control, bit 12 of the I_2 field, controls whether the transaction is allowed to modify an access register. The effective allow AR modification control is the logical AND of the A control in the TBEGIN instruction for the current nesting level and for all outer levels.

If the effective A control is zero, the transaction will be aborted with abort code 11 (restricted instruction) if an attempt is made to modify any access register. If the effective A control is one, the transaction will not be aborted if an access register is modified (absent of any other abort condition).

Allow Floating Point Operation (F) **214**: The F control, bit 13 of the I_2 field, controls whether the transaction is allowed to execute specified floating point instructions. The effective allow floating point operation control is the logical AND of the F control in the TBEGIN instruction for the current nesting level and for all outer levels.

If the effective F control is zero, then (a) the transaction will be aborted with abort code 11 (restricted instruction) if an attempt is made to execute a floating point instruction, and (b) the data exception code (DXC) in byte 2 of the floating point control register (FPCR) will not be set by any data exception program exception condition. If the effective F control is one, then (a) the transaction will not be aborted if an attempt is made to execute a floating point instruction (absent any other abort condition), and (b) the DXC in the FPCR may be set by a data exception program exception condition.

Program Interruption Filtering Control (PIFC) **216**: Bits 14-15 of the I_2 field are the program interruption filtering control (PIFC). The PIFC controls whether certain classes of program exception conditions (e.g., addressing exception, data exception, operation exception, protection exception, etc.) that occur while the CPU is in the transactional execution mode result in an interruption.

The effective PIFC is the highest value of the PIFC in the TBEGIN instruction for the current nesting level and for all outer levels. When the effective PIFC is zero, all program exception conditions result in an interruption. When the effective PIFC is one, program exception conditions having a transactional execution class of 1 and 2 result in an interruption. (Each program exception condition is assigned at least one transactional execution class, depending on the severity of the exception. Severity is based on the likelihood of recovery during a repeated execution of the transactional execution, and whether the operating system needs to see the interruption.) When the effective PIFC is two, program exception conditions having a transactional execution class of 1 result in an interruption. A PIFC of 3 is reserved.

Bits 8-11 of the I_2 field (bits 40-43 of the instruction) are reserved and should contain zeros; otherwise, the program may not operate compatibly in the future.

One embodiment of a format of a Transaction Begin constrained (TBEGINC) instruction is described with reference to FIGS. 3A-3B. In one example, TBEGINC **300** (FIG. 3A) includes an opcode field **302** that includes an opcode specifying a transaction begin constrained operation; a base field (B_1) **304**; a displacement field (D_1) **306**; and an immediate field (I_2) **308**. The contents of the general register specified by B_1 **304** are added to D_1 **306** to obtain the first operand address. However, with the transaction begin constrained instruction, the first operand address is not used to access storage. Instead, the B_1 field of the instruction includes zeros; otherwise, a specification exception is recognized.

In one embodiment, the I_2 field includes various controls, an example of which is depicted in FIG. 3B.

The bits of the I_2 field are defined as follows, in one example:

General Register Save Mask (GRSM) 310: Bits 0-7 of the I_2 field contain the general register save mask (GRSM). Each bit of the GRSM represents an even-odd pair of general registers, where bit 0 represents registers 0 and 1, bit 1 represents registers 2 and 3, and so forth. When a bit in the GRSM is zero, the corresponding register pair is not saved. When a bit in the GRSM is one, the corresponding register pair is saved in a model-dependent location that is not directly accessible by the program.

If the transaction aborts, saved register pairs are restored to their contents when the outermost TRANSACTION BEGIN instruction was executed. The contents of all other (unsaved) general registers are not restored when a constrained transaction aborts.

When TBEGINC is used to continue execution in the nonconstrained transaction execution mode, the general register save mask is ignored.

Allow AR Modification (A) 312: The A control, bit 12 of the I_2 field, controls whether the transaction is allowed to modify an access register. The effective allow-AR-modification control is the logical AND of the A control in the TBEGINC instruction for the current nesting level and for any outer TBEGIN or TBEGINC instructions.

If the effective A control is zero, the transaction will be aborted with abort code 11 (restricted instruction) if an attempt is made to modify any access register. If the effective A control is one, the transaction will not be aborted if an access register is modified (absent of any other abort condition).

Bits 8-11 and 13-15 of the I_2 field (bits 40-43 and 45-47 of the instruction) are reserved and should contain zeros.

The end of a Transaction Begin instruction is specified, in one example, by a TRANSACTION END (TEND) instruction, a format of which is depicted in FIG. 4. As one example, a TEND instruction **400** includes an opcode field **402** that includes an opcode specifying a transaction end operation.

A number of terms are used with respect to the transactional execution facility, and therefore, solely for convenience, a list of terms is provided below in alphabetical order. In one embodiment, these terms have the following definition:

Abort: A transaction aborts when it is ended prior to a TRANSACTION END instruction that results in a transaction nesting depth of zero, or when a CONDITIONAL

TRANSACTION END instruction sets a condition code of zero. When a transaction aborts, the following occurs, in one embodiment:

Transactional store accesses made by any and all levels of the transaction are discarded (that is, not committed). Non-transactional store accesses made by any and all levels of the transaction are committed.

Registers designated by the general register save mask (GRSM) of the outermost TRANSACTION BEGIN instruction are restored to their contents prior to the transactional execution (that is, to their contents at execution of the outermost TRANSACTION BEGIN instruction). General registers not designated by the general register save mask of the outermost TRANSACTION BEGIN instruction are not restored.

Access registers, floating-point registers, and the floating-point control register are not restored. Any changes made to these registers during transaction execution are retained when the transaction aborts.

A transaction may be aborted due to a variety of reasons, including attempted execution of a restricted instruction, attempted modification of a restricted resource, transactional conflict, exceeding various CPU resources, any interpretive-execution interception condition, any interruption, a TRANSACTION ABORT instruction, and other reasons. A transaction-abort code provides specific reasons why a transaction may be aborted.

One example of a format of a TRANSACTION ABORT (TABORT) instruction is described with reference to FIG. 5. As one example, a TABORT instruction **500** includes an opcode field **502** that includes an opcode specifying a transaction abort operation; a base field (B_2) **504**; and a displacement field (D_2) **506**. When the B_2 field is nonzero, the contents of the general register specified by B_2 **504** are added to D_2 **506** to obtain a second operand address; otherwise, the second operand address is formed solely from the D_2 field, and the B_2 field is ignored. The second operand address is not used to address data; instead, the address forms the transaction abort code which is placed in a transaction diagnostic block during abort processing. Address computation for the second operand address follows the rules of address arithmetic: in the 24-bit addressing mode, bits 0-29 are set to zeros; in the 31-bit addressing mode, bits 0-32 are set to zeros.

Commit: At the completion of an outermost TRANSACTION END instruction, or at the completion of a CONDITIONAL TRANSACTION END instruction that sets condition code 0, the CPU commits the store accesses made by the transaction (i.e., the outermost transaction and any nested levels) such that they are visible to other CPUs and the I/O subsystem. As observed by other CPUs and by the I/O subsystem, all fetch and store accesses made by all nested levels of the transaction appear to occur as a single concurrent operation when the commit occurs.

The contents of the general registers, access registers, floating-point registers, and the floating-point control register are not modified by the commit process. Any changes made to these registers during transactional execution are retained when the transaction's stores are committed.

Conflict: A transactional access made by one CPU conflicts with either (a) a transactional access or non-transactional access made by another CPU, or (b) the non-transactional access made by the I/O subsystem, if both accesses are to any location within the same cache line, and one or more of the accesses is a store.

A conflict may be detected by a CPU's speculative execution of instructions, even though the conflict may not be detected in the conceptual sequence.

Constrained Transaction: A constrained transaction is a transaction that executes in the constrained transactional execution mode and is subject to the following limitations:

A subset of the general instructions is available.

A limited number of instructions may be executed.

A limited number of storage-operand locations may be accessed.

The transaction is limited to a single nesting level.

In the absence of repeated interruptions or conflicts with other CPUs or the I/O subsystem, a constrained transaction eventually completes, thus an abort-handler routine is not required.

When a TRANSACTION BEGIN constrained (TBEGINC) instruction is executed while the CPU is already in the nonconstrained transaction execution mode, execution continues as a nested nonconstrained transaction.

Constrained Transactional Execution Mode: When the transaction nesting depth is zero, and a transaction is initiated by a TBEGINC instruction, the CPU enters the constrained transactional execution mode. While the CPU is in the constrained transactional execution mode, the transaction nesting depth is one.

Nested Transaction: When the TRANSACTION BEGIN instruction is issued while the CPU is in the nonconstrained transactional execution mode, the transaction is nested.

The transactional execution facility uses a model called flattened nesting. In the flattened nesting mode, stores made by an inner transaction are not observable by other CPUs and by the I/O subsystem until the outermost transaction commits its stores. Similarly, if a transaction aborts, all nested transactions abort, and all transactional stores of all nested transactions are discarded.

One example of nested transactions is depicted in FIG. 6. As shown, a first TBEGIN 600 starts an outermost transaction 601, TBEGIN 602 starts a first nested transaction, and TBEGIN 604 starts a second nested transaction. In this example, TBEGIN 604 and TEND 606 define an innermost transaction 608. When TEND 610 executes, transactional stores are committed 612 for the outermost transaction and all inner transactions.

Nonconstrained Transaction: A nonconstrained transaction is a transaction that executes in the nonconstrained transactional execution mode. Although a nonconstrained transaction is not limited in the manner as a constrained transaction, it may still be aborted due to a variety of causes.

Nonconstrained Transactional Execution Mode: When a transaction is initiated by the TBEGIN instruction, the CPU enters the nonconstrained transactional execution mode. While the CPU is in the nonconstrained transactional execution mode, the transaction nesting depth may vary from one to the maximum transaction nesting depth.

Non-Transactional Access: Non-transactional accesses are storage operand accesses made by the CPU when it is not in the transactional execution mode (that is, classic storage accesses outside of a transaction). Further, accesses made by the I/O subsystem are non-transactional accesses. Additionally, the NONTRANSACTIONAL STORE instruction may be used to cause a non-transactional store access while the CPU is in the nonconstrained transactional execution mode.

Outer/Outermost Transaction: A transaction with a lower-numbered transaction nesting depth is an outer transaction. A transaction with a transaction nesting depth value of one is the outermost transaction.

An outermost TRANSACTION BEGIN instruction is one that is executed when the transaction nesting depth is initially zero. An outermost TRANSACTION END instruction is one that causes the transaction nesting depth to transition from one to zero. Further, a CONDITIONAL TRANSACTION END instruction that sets the condition code to zero may also be considered to be the outermost form of the instruction. A constrained transaction is the outermost transaction, in this embodiment.

Program Interruption Filtering: When a transaction is aborted due to certain program exception conditions, the program can optionally prevent the interruption from occurring. This technique is called program-interruption filtering. Program interruption filtering is subject to the transactional class of the interruption, the effective program interruption filtering control from the TRANSACTION BEGIN instruction, and the transactional execution program interruption filtering override in control register 0.

Transaction: A transaction includes the storage-operand accesses made, and selected general registers altered, while the CPU is in the transaction execution mode. For a non-constrained transaction, storage-operand accesses may include both transactional accesses and non-transactional accesses. For a constrained transaction, storage-operand accesses are limited to transactional accesses. As observed by other CPUs and by the I/O subsystem, all storage-operand accesses made by the CPU while in the transaction execution mode appear to occur as a single concurrent operation. If a transaction is aborted, transactional store accesses are discarded, and any registers designated by the general register save mask of the outermost TRANSACTION BEGIN instruction are restored to their contents prior to transactional execution.

Transactional Accesses: Transactional accesses are storage operand accesses made while the CPU is in the transactional execution mode, with the exception of accesses made by the NONTRANSACTIONAL STORE instruction.

Transactional Execution Mode: The term transactional execution mode (a.k.a., transaction execution mode) describes the common operation of both the nonconstrained and the constrained transactional execution modes. Thus, when the operation is described, the terms nonconstrained and constrained are used to qualify the transactional execution mode.

When the transaction nesting depth is zero, the CPU is not in the transactional execution mode (also called the non-transactional execution mode).

As observed by the CPU, fetches and stores made in the transactional execution mode are no different than those made while not in the transactional execution mode.

In one embodiment of the z/Architecture, the transactional execution facility is under the control of bits 8-9 of control register 0, bits 61-63 of control register 2, the transaction nesting depth, the transaction diagnostic block address, and the transaction abort program status word (PSW).

Following an initial CPU reset, the contents of bit positions 8-9 of control register 0, bit positions 62-63 of control register 2, and the transaction nesting depth are set to zero. When the transactional execution control, bit 8 of control register 0, is zero, the CPU cannot be placed into the transactional execution mode.

Further details regarding the various controls are described below.

As indicated, the transactional execution facility is controlled by two bits in control register zero and three bits in control register two. For instance:

Control Register 0 Bits: The bit assignments are as follows, in one embodiment:

Transactional Execution Control (TXC): Bit 8 of control register zero is the transactional execution control. This bit provides a mechanism whereby the control program (e.g., operating system) can indicate whether or not the transactional execution facility is usable by the program. Bit 8 is to be one to successfully enter the transactional execution mode.

When bit 8 of control register 0 is zero, attempted execution of the CONDITIONAL TRANSACTION END, EXTRACT TRANSACTION NESTING DEPTH, TRANSACTION BEGIN and TRANSACTION END instructions results in a special operation execution.

Transaction Execution Program Interruption Filtering Override (PIFO): Bit 9 of control register zero is the transactional execution program interruption filtering override. This bit provides a mechanism by which the control program can ensure that any program exception condition that occurs while the CPU is in the transactional execution mode results in an interruption, regardless of the effective program interruption filtering control specified or implied by the TRANSACTION BEGIN instruction(s).

Control Register 2 Bits: The assignments are as follows, in one embodiment:

Transaction Diagnostic Scope (TDS): Bit 61 of control register 2 controls the applicability of the transaction diagnosis control (TDC) in bits 62-63 of the register, as follows:

TDS

Value	Meaning
0	The TDC applies regardless of whether the CPU is in the problem or supervisor state.
1	The TDC applies only when the CPU is in the problem state. When the CPU is in the supervisor state, processing is as if the TDC contained zero.

Transaction Diagnostic Control (TDC): Bits 62-63 of control register 2 are a 2-bit unsigned integer that may be used to cause transactions to be randomly aborted for diagnostic purposes. The encoding of the TDC is as follows, in one example:

TDC

Value	Meaning
0	Normal operation; transactions are not aborted as a result of the TDC.
1	Abort every transaction at a random instruction, but before execution of the outermost TRANSACTION END instruction or CONDITIONAL TRANSACTION END instruction that sets condition code zero.
2	Abort random transactions at a random instruction.
3	Reserved

When a transaction is aborted due to a nonzero TDC, then either of the following may occur:

The abort code is set to any of the codes 7-11, 13-18, or 255, with the value of the code randomly chosen by the CPU; the condition code is set corresponding to the abort code.

For a nonconstrained transaction, the condition code is set to one. In this case, the abort code is not applicable.

It is model dependent whether TDC value 1 is implemented. If not implemented, a value of 1 acts as if 2 was specified.

For a constrained transaction, a TDC value of 1 is treated as if a TDC value of 2 was specified.

If a TDC value of 3 is specified, the results are unpredictable.

Transaction Diagnostic Block Address (TDBA)

A valid transaction diagnostic block address (TDBA) is set from the first operand address of the outermost TRANSACTION BEGIN (TBEGIN) instruction when the B_1 field of the instruction is nonzero. When the CPU is in the primary space or access register mode, the TDBA designates a location in the primary address space. When the CPU is in the secondary space, or home space mode, the TDBA designates a location in the secondary or home address space, respectively. When DAT (Dynamic Address Translation) is off, the TDBA designates a location in real storage.

The TDBA is used by the CPU to locate the transaction diagnostic block—called the TBEGIN-specified TDB—if the transaction is subsequently aborted. The rightmost three bits of the TDBA are zero, meaning that the TBEGIN-specified TDB is on a doubleword boundary.

When the B_1 field of an outermost TRANSACTION BEGIN (TBEGIN) instruction is zero, the transactional diagnostic block address is invalid, and no TBEGIN-specified TDB is stored if the transaction is subsequently aborted.

Transaction Abort PSW (TAPSW)

During execution of the TRANSACTION BEGIN (TBEGIN) instruction when the nesting depth is initially zero, the transaction abort PSW is set to the contents of the current PSW; and the instruction address of the transaction abort PSW designates the next sequential instruction (that is, the instruction following the outermost TBEGIN). During execution of the TRANSACTION BEGIN constrained (TBEGINC) instruction when the nesting depth is initially zero, the transaction abort PSW is set to the contents of the current PSW, except that the instruction address of the transaction abort PSW designates the TBEGINC instruction (rather than the next sequential instruction following the TBEGINC).

When a transaction is aborted, the condition code in the transaction abort PSW is replaced with a code indicating the severity of the abort condition. Subsequently, if the transaction was aborted due to causes that do not result in an interruption, the PSW is loaded from the transaction abort PSW; if the transaction was aborted due to causes that result in an interruption, the transaction abort PSW is stored as the interruption old PSW.

The transaction abort PSW is not altered during the execution of any inner TRANSACTION BEGIN instruction.

Transaction Nesting Depth (TND)

The transaction nesting depth is, for instance, a 16-bit unsigned value that is incremented each time a TRANSACTION BEGIN instruction is completed with condition code 0 and decremented each time a TRANSACTION END or CONDITIONAL TRANSACTION END instruction is completed with condition code zero. The transaction nesting depth is reset to zero when a transaction is aborted or by CPU reset.

In one embodiment, a maximum TND of 15 is implemented.

In one implementation, when the CPU is in the constrained transactional execution mode, the transaction nesting depth is one. Additionally, although the maximum TND

can be represented as a 4-bit value, the TND is defined to be a 16-bit value to facilitate its inspection in the transaction diagnostic block.

Transaction Diagnostic Block (TDB)

When a transaction is aborted, various status information may be saved in a transaction diagnostic block (TDB), as follows:

1. TBEGIN-specified TDB: For a nonconstrained transaction, when the B_1 field of the outermost TBEGIN instruction is nonzero, the first operand address of the instruction designates the TBEGIN-specified TDB. This is an application program specified location that may be examined by the application's abort handler.

2. Program-Interruption (PI) TDB: If a nonconstrained transaction is aborted due to a non-filtered program exception condition, or if a constrained transaction is aborted due to any program exception condition (that is, any condition that results in a program interruption being recognized), the PI-TDB is stored into locations in the prefix area. This is available for the operating system to inspect and log out in any diagnostic reporting that it may provide.

3. Interception TDB: If the transaction is aborted due to any program exception condition that results in interception (that is, the condition causes interpretive execution to end and control to return to the host program), a TDB is stored into a location specified in the state description block for the guest operating system.

The TBEGIN-specified TDB is only stored, in one embodiment, when the TDB address is valid (that is, when the outermost TBEGIN instruction's B_1 field is nonzero).

For aborts due to unfiltered program exception conditions, only one of either the PI-TDB or Interception TDB will be stored. Thus, there may be zero, one, or two TDBs stored for an abort.

Further details regarding one example of each of the TDBs are described below:

TBEGIN-specified TDB: The 256-byte location specified by a valid transaction diagnostic block address. When the transaction diagnostic block address is valid, the TBEGIN-specified TDB is stored on a transaction abort. The TBEGIN-specified TDB is subject to all storage protection mechanisms that are in effect at the execution of the outermost TRANSACTION BEGIN instruction. A PER (Program Event Recording) storage alteration event for any portion of the TBEGIN-specified TDB is detected during the execution of the outermost TBEGIN, not during the transaction abort processing.

One purpose of PER is to assist in debugging programs. It permits the program to be alerted to the following types of events, as examples:

Execution of a successful branch instruction. The option is provided of having an event occur only when the branch target location is within the designated storage area.

Fetching of an instruction from the designated storage area.

Alteration of the contents of the designated storage area. The option is provided of having an event occur only when the storage area is within designated address spaces.

Execution of a STORE USING REAL ADDRESS instruction.

Execution of the CONDITIONAL TRANSACTION END or TRANSACTION END instruction.

The program can selectively specify that one or more of the above types of events be recognized, except that the

event for STORE USING REAL ADDRESS can be specified only along with the storage alteration event. The information concerning a PER event is provided to the program by means of a program interruption, with the cause of the interruption being identified in the interruption code.

When the transaction diagnostic block address is not valid, a TBEGIN-specified TDB is not stored.

Program-Interruption TDB: Real locations 6,144-6,399 (1800-18FF hex). The program interruption TDB is stored when a transaction is aborted due to program interruption. When a transaction is aborted due to other causes, the contents of the program interruption TDB are unpredictable.

The program interruption TDB is not subject to any protection mechanism. PER storage alteration events are not detected for the program interruption TDB when it is stored during a program interruption.

Interception TDB: The 256-byte host real location specified by locations **488-495** of the state description. The interception TDB is stored when an aborted transaction results in a guest program interruption interception (that is, interception code 8). When a transaction is aborted due to other causes, the contents of the interception TDB are unpredictable. The interception TDB is not subject to any protection mechanism.

As depicted in FIG. 7, the fields of a transaction diagnostic block **700** are as follows, in one embodiment:

Format **702**: Byte 0 contains a validity and format indication, as follows:

Value	Meaning
0	The remaining fields of the TDB are unpredictable.
1	A format-1 TDB, the remaining fields of which are described below.
2-255	Reserved

A TDB in which the format field is zero is referred to as a null TDB.

Flags **704**: Byte 1 contains various indications, as follows:

Conflict Token Validity (CTV): When a transaction is aborted due to a fetch or store conflict (that is, abort codes 9 or 10, respectively), bit 0 of byte 1 is the conflict token validity indication. When the CTV indication is one, the conflict token **710** in bytes 16-23 of the TDB contain the logical address at which the conflict was detected. When the CTV indication is zero, bytes 16-23 of the TDB are unpredictable.

When a transaction is aborted due to any reason other than a fetch or store conflict, bit 0 of byte 1 is stored as zero.

Constrained-Transaction Indication (CTI): When the CPU is in the constrained transactional execution mode, bit 1 of byte 1 is set to one. When the CPU is in the nonconstrained transactional execution mode, bit 1 of byte 1 is set to zero.

Reserved: Bits 2-7 of byte 1 are reserved, and stored as zeros.

Transaction Nesting Depth (TND) **706**: Bytes 6-7 contain the transaction nesting depth when the transaction was aborted.

Transaction Abort Code (TAC) **708**: Bytes 8-15 contain a 64-bit unsigned transaction abort code. Each code indicates a reason for a transaction being aborted.

It is model dependent whether the transaction abort code is stored in the program interruption TDB when a transaction is aborted due to conditions other than a program interruption.

Conflict Token **710**: For transactions that are aborted due to fetch or store conflict (that is, abort codes 9 and 10, respectively), bytes 16-23 contain the logical address of the storage location at which the conflict was detected. The conflict token is meaningful when the CTV bit, bit 0 of byte 1, is one.

When the CTV bit is zero, bytes 16-23 are unpredictable.

Because of speculative execution by the CPU, the conflict token may designate a storage location that would not necessarily be accessed by the transaction's conceptual execution sequence.

Aborted Transaction Instruction Address (ATIA) **712**: Bytes 24-31 contain an instruction address that identifies the instruction that was executing when an abort was detected. When a transaction is aborted due to abort codes 2, 5, 6, 11, 13, 17, 18, or 256 or higher, or when a transaction is aborted due to abort codes 4 or 12 and the program exception condition is nullifying, the ATIA points directly to the instruction that was being executed. When a transaction is aborted due to abort codes 4 or 12, and the program exception condition is not nullifying, the ATIA points past the instruction that was being executed.

When a transaction is aborted due to abort codes 7-10, 14-16, or 255, the ATIA does not necessarily indicate the exact instruction causing the abort, but may point to an earlier or later instruction within the transaction.

If a transaction is aborted due to an instruction that is the target of an execute-type instruction, the ATIA identifies the execute-type instruction, either pointing to the instruction or past it, depending on the abort code. The ATIA does not indicate the target of the execute-type instruction.

The ATIA is subject to the addressing mode when the transaction is aborted. In the 24-bit addressing mode, bits 0-40 of the field contain zeros. In the 31-bit addressing mode, bits 0-32 of the field contain zeros.

It is model dependent whether the aborted transaction instruction address is stored in the program interruption TDB when a transaction is aborted due to conditions other than a program interruption.

When a transaction is aborted due to abort code 4 or 12, and the program exception condition is not nullifying, the ATIA does not point to the instruction causing the abort. By subtracting the number of halfwords indicated by the interruption length code (ILC) from the ATIA, the instruction causing the abort can be identified in conditions that are suppressing or terminating, or for non-PER events that are completing. When a transaction is aborted due to a PER event, and no other program exception condition is present, the ATIA is unpredictable.

When the transaction diagnostic block address is valid, the ILC may be examined in program interruption identification (PhD) in bytes 36-39 of the TBEGIN-specified TDB. When filtering does not apply, the ILC may be examined in the PhD at location **140-143** in real storage.

Exception Access Identification (EAID) **714**: For transactions that are aborted due to certain filtered program exception conditions, byte 32 of the TBEGIN-specified TDB contains the exception access identification. In one example of the z/Architecture, the format of the EAID, and the cases for which it is stored, are the same as those described in real location **160** when the exception condition results in an interruption, as described in the above-incorporated by reference Principles of Operation.

For transactions that are aborted for other reasons, including any exception conditions that result in a program interruption, byte 32 is unpredictable. Byte 32 is unpredictable in the program interruption TDB.

This field is stored only in the TDB designated by the transaction diagnostic block address; otherwise, the field is reserved. The EAID is stored only for access list controlled or DAT protection, ASCE-type, page translation, region first translation, region second translation, region third translation, and segment translation program exception conditions.

Data Exception Code (DXC) **716**: For transactions that are aborted due to filtered data exception program exception conditions, byte 33 of the TBEGIN specified TDB contains the data exception code. In one example of the z/Architecture, the format of the DXC, and the cases for which it is stored, are the same as those described in real location **147** when the exception condition results in an interruption, as described in the above-incorporated by reference Principles of Operation. In one example, location **147** includes the DXC.

For transactions that are aborted for other reasons, including any exception conditions that result in a program interruption, byte 33 is unpredictable. Byte 33 is unpredictable in the program interruption TDB.

This field is stored only in the TDB designated by the transaction diagnostic block address; otherwise, the field is reserved. The DXC is stored only for data program exception conditions.

Program Interruption Identification (PhD) **718**: For transactions that are aborted due to filtered program exception conditions, bytes 36-39 of the TBEGIN-specified TDB contain the program interruption identification. In one example of the z/Architecture, the format of the PIID is the same as that described in real locations **140-143** when the condition results in an interruption (as described in the above-incorporated by reference Principles of Operation), except that the instruction length code in bits 13-14 of the PIID is respective to the instruction at which the exception condition was detected.

For transactions that are aborted for other reasons, including exception conditions that result in a program interruption, bytes 36-39 are unpredictable. Bytes 36-39 are unpredictable in the program interruption TDB.

This field is stored only in the TDB designated by the transaction diagnostic block address; otherwise, the field is reserved. The program interruption identification is only stored for program exception conditions.

Translation Exception Identification (TED) **720**: For transactions that are aborted due to any of the following filtered program exception conditions, bytes 40-47 of the TBEGIN-specified TDB contain the translation exception identification.

- Access list controlled or DAT protection
- ASCE-type
- Page translation
- Region-first translation
- Region-second translation
- Region-third translation
- Segment translation exception

In one example of the z/Architecture, the format of the TED is the same as that described in real locations **168-175** when the condition results in an interruption, as described in the above-incorporated by reference Principles of Operation.

For transactions that are aborted for other reasons, including exception conditions that result in a program interruption, bytes 40-47 are unpredictable. Bytes 40-47 are unpredictable in the program interruption TDB.

This field is stored only in the TDB designated by the transaction diagnostic block address; otherwise, the field is reserved.

23

Breaking Event Address **722**: For transactions that are aborted due to filtered program exception conditions, bytes 48-55 of the TBEGIN-specified TDB contain the breaking event address. In one example of the z/Architecture, the format of the breaking event address is the same as that described in real locations **272-279** when the condition results in an interruption, as described in the above-incorporated by reference Principles of Operation.

For transactions that are aborted for other reasons, including exception conditions that result in a program interruption, bytes 48-55 are unpredictable. Bytes 48-55 are unpredictable in the program interruption TDB.

This field is stored only in the TDB designated by the transaction diagnostic block address; otherwise, the field is reserved.

Further details relating to breaking events are described below.

In one embodiment of the z/Architecture, when the PER-3 facility is installed, it provides the program with the address of the last instruction to cause a break in the sequential execution of the CPU. Breaking event address recording can be used as a debugging assist for wild branch detection. This facility provides, for instance, a 64-bit register in the CPU, called the breaking event address register. Each time an instruction other than TRANSACTION ABORT causes a break in the sequential instruction execution (that is, the instruction address in the PSW is replaced, rather than incremented by the length of the instruction), the address of that instruction is placed in the breaking event address register. Whenever a program interruption occurs, whether or not PER is indicated, the current contents of the breaking event address register are placed in real storage locations **272-279**.

If the instruction causing the breaking event is the target of an execute-type instruction (EXECUTE or EXECUTE RELATIVE LONG), then the instruction address used to fetch the execute-type instruction is placed in the breaking event address register.

In one embodiment of the z/Architecture, a breaking event is considered to occur whenever one of the following instructions causes branching: BRANCH AND LINK (BAL, BALR); BRANCH AND SAVE (BAS, BASR); BRANCH AND SAVE AND SET MODE (BASSM); BRANCH AND SET MODE (BSM); BRANCH AND STACK (BAKR); BRANCH ON CONDITION (BC, BCR); BRANCH ON COUNT (BCT, BCTR, BCTG, BCTGR); BRANCH ON INDEX HIGH (BXH, BXHG); BRANCH ON INDEX LOW OR EQUAL (BXLE, BXLEG); BRANCH RELATIVE ON CONDITION (BRC); BRANCH RELATIVE ON CONDITION LONG (BRCL); BRANCH RELATIVE ON COUNT (BRCT, BRCTG); BRANCH RELATIVE ON INDEX HIGH (BRXH, BRXHG); BRANCH RELATIVE ON INDEX LOW OR EQUAL (BRXLE, BRXLG); COMPARE AND BRANCH (CRB, CGRB); COMPARE AND BRANCH RELATIVE (CRJ, CGRJ); COMPARE IMMEDIATE AND BRANCH (CM, CGIB); COMPARE IMMEDIATE AND BRANCH RELATIVE (CIJ, CGIJ); COMPARE LOGICAL AND BRANCH (CLRB, CLGRB); COMPARE LOGICAL AND BRANCH RELATIVE (CLRJ, CLGRJ); COMPARE LOGICAL IMMEDIATE AND BRANCH (CLIB, CLUB); and COMPARE LOGICAL IMMEDIATE AND BRANCH RELATIVE (CLIJ, CLGIJ).

A breaking event is also considered to occur whenever one of the following instructions completes: BRANCH AND SET AUTHORITY (BSA); BRANCH IN SUBSPACE GROUP (BSG); BRANCH RELATIVE AND SAVE;

24

BRANCH RELATIVE AND SAVE LONG (BRASL); LOAD PSW (LPSW); LOAD PSW EXTENDED (LPSWE); PROGRAM CALL (PC); PROGRAM RETURN (PR); PROGRAM TRANSFER (PT); PROGRAM TRANSFER WITH INSTANCE (PTI); RESUME PROGRAM (RP); and TRAP (TRAP2, TRAP4).

A breaking event is not considered to occur as a result of a transaction being aborted (either implicitly or as a result of the TRANSACTION ABORT instruction).

Model Dependent Diagnostic Information **724**: Bytes 112-127 contain model dependent diagnostic information.

For all abort codes except 12 (filtered program interruption), the model dependent diagnostic information is saved in each TDB that is stored.

In one embodiment, the model dependent diagnostic information includes the following:

Bytes 112-119 contain a vector of 64 bits called the transactional execution branch indications (TXBI). Each of the first 63 bits of the vector indicates the results of executing a branching instruction while the CPU was in the transactional execution mode, as follows:

Value	Meaning
0	The instruction completed without branching.
1	The instruction completed with branching.

Bit 0 represents the result of the first such branching instruction; bit 1 represents the result of the second such instruction, and so forth.

If fewer than 63 branching instructions were executed while the CPU was in the transactional execution mode, the rightmost bits that do not correspond to branching instructions are set to zeros (including bit 63). When more than 63 branching instructions were executed, bit 63 of the TXBI is set to one.

Bits in the TXBI are set by instructions which are capable of causing a breaking event, as listed above, except for the following:

Any restricted instruction does not cause a bit to be set in the TXBI.

For instructions of, for instance, the z/Architecture, when the M_1 field of the BRANCH ON CONDITION, BRANCH RELATIVE ON CONDITION, or BRANCH RELATIVE ON CONDITION LONG instruction is zero, or when the R_2 field of the following instructions is zero, it is model dependent whether the execution of the instruction causes a bit to be set in the TXBI.

BRANCH AND LINK (BALR); BRANCH AND SAVE (BASR); BRANCH AND SAVE AND SET MODE (BASSM); BRANCH AND SET MODE (BSM); BRANCH ON CONDITION (BCR); and BRANCH ON COUNT (BCTR, BCTGR)

For abort conditions that were caused by a host access exception, bit position 0 of byte 127 is set to one. For all other abort conditions, bit position 0 of byte 127 is set to zero.

For abort conditions that were detected by the load/store unit (LSU), the rightmost five bits of byte 127 contain an indication of the cause. For abort conditions that were not detected by the LSU, byte 127 is reserved.

General Registers **730**: Bytes 128-255 contain the contents of general registers 0-15 at the time the transaction was aborted. The registers are stored in ascending order, begin-

ning with general register 0 in bytes 128-135, general register 1 in bytes 136-143, and so forth.

Reserved: All other fields are reserved. Unless indicated otherwise, the contents of reserved fields are unpredictable.

As observed by other CPUs and the I/O subsystem, storing of the TDB(s) during a transaction abort is a multiple access reference occurring after any non-transactional stores.

A transaction may be aborted due to causes that are outside the scope of the immediate configuration in which it executes. For example, transient events recognized by a hypervisor (such as LPAR or z/VM) may cause a transaction to be aborted.

The information provided in the transaction diagnostic block is intended for diagnostic purposes and is substantially correct. However, because an abort may have been caused by an event outside the scope of the immediate configuration, information such as the abort code or program interruption identification may not accurately reflect conditions within the configuration, and thus, should not be used in determining program action.

In addition to the diagnostic information saved in the TDB, when a transaction is aborted due to any data exception program exception condition and both the AFP register control, bit 45 of control register 0, and the effective allow floating point operation control (F) are one, the data exception code (DXC) is placed into byte 2 of the floating point control register (FPCR), regardless of whether filtering applies to the program exception condition. When a transaction is aborted, and either or both the AFP register control or effective allow floating point operation control are zero, the DXC is not placed into the FPCR.

In one embodiment, as indicated herein, when the transaction execution facility is installed, a number of general instructions are provided, including, for instance, EXTRACT TRANSACTION NESTING DEPTH, NON-TRANSACTIONAL STORE, TRANSACTION ABORT, TRANSACTION BEGIN and TRANSACTION END.

Further, when the conditional transaction end facility is installed, the CONDITIONAL TRANSACTION END instruction is provided.

When the CPU is in the transaction execution mode, attempted execution of certain instructions is restricted and causes the transaction to be aborted.

When issued in the constrained transactional execution mode, attempted execution of restricted instructions may also result in a transaction constraint program interruption, or may result in execution proceeding as if the transaction was not constrained.

In one example of the z/Architecture, restricted instructions include, as examples, the following non-privileged instructions: COMPARE AND SWAP AND STORE; MODIFY RUNTIME INSTRUMENTATION CONTROLS; PERFORM LOCKED OPERATION; PREFETCH DATA (RELATIVE LONG), when the code in the M_1 field is 6 or 7; STORE CHARACTERS UNDER MASK HIGH, when the M_3 field is zero and the code in the R_1 field is 6 or 7; STORE FACILITY LIST EXTENDED; STORE RUNTIME INSTRUMENTATION CONTROLS; SUPERVISOR CALL; and TEST RUNTIME INSTRUMENTATION CONTROLS.

In the above list, COMPARE AND SWAP AND STORE and PERFORM LOCKED OPERATION are complex instructions which can be more efficiently implemented by the use of basic instructions in the TX mode. The cases for PREFETCH DATA and PREFETCH DATA RELATIVE LONG are restricted as the codes of 6 and 7 release a cache

line, necessitating the commitment of the data potentially prior to the completion of a transaction. SUPERVISOR CALL is restricted as it causes an interruption (which causes a transaction to be aborted).

Under the conditions listed below, the following instructions are restricted:

BRANCH AND LINK (BALR), BRANCH AND SAVE (BASR), and BRANCH AND SAVE AND SET MODE, when the R_2 field of the instruction is nonzero and branch tracing is enabled.

BRANCH AND SAVE AND SET MODE and BRANCH AND SET MODE, when the R_2 field is nonzero and mode tracing is enabled; SET ADDRESSING MODE, when mode tracing is enabled.

CONDITIONAL TRANSACTION END when the transaction nesting depth is greater than one.

MONITOR CALL, when a monitor event condition is recognized.

The above list includes instructions that may form trace entries. If these instructions were allowed to execute transactionally and formed trace entries, and the transaction subsequently aborted, the trace table pointer in control register 12 would be advanced, but the stores to the trace table would be discarded. This would leave an inconsistent gap in the trace table; thus, the instructions are restricted in the cases where they would form trace entries.

When the CPU is in the transactional execution mode, it is model dependent whether the following instructions are restricted: CIPHER MESSAGE; CIPHER MESSAGE WITH CFB; CIPHER MESSAGE WITH CHAINING; CIPHER MESSAGE WITH COUNTER; CIPHER MESSAGE WITH OFB; COMPRESSION CALL; COMPUTE INTERMEDIATE MESSAGE DIGEST; COMPUTE LAST MESSAGE DIGEST; COMPUTE MESSAGE AUTHENTICATION CODE; CONVERT UNICODE-16 TO UNICODE-32; CONVERT UNICODE-16 TO UNICODE-8; CONVERT UNICODE-32 TO UNICODE-16; CONVERT UNICODE-32 TO UNICODE-8; CONVERT UNICODE-8 TO UNICODE-16; CONVERT UNICODE-8 TO UNICODE-32; PERFORM CRYPTOGRAPHIC COMPUTATION; RUNTIME INSTRUMENTATION OFF; and RUNTIME INSTRUMENTATION ON.

Each of the above instructions is either currently implemented by the hardware co-processor, or has been in past machines, and thus, is considered restricted.

When the effective allow AR modification (A) control is zero, the following instructions are restricted: COPY ACCESS; LOAD ACCESS MULTIPLE; LOAD ADDRESS EXTENDED; and SET ACCESS.

Each of the above instructions causes the contents of an access register to be modified. If the A control in the TRANSACTION BEGIN instruction is zero, then the program has explicitly indicated that access register modification is not to be allowed.

When the effective allow floating point operation (F) control is zero, floating point instructions are restricted.

Under certain circumstances, the following instructions may be restricted: EXTRACT CPU TIME; EXTRACT PSW; STORE CLOCK; STORE CLOCK EXTENDED; and STORE CLOCK FAST.

Each of the above instructions is subject to an interception control in the interpretative execution state description. If the hypervisor has set the interception control for these instructions, then their execution may be prolonged due to hypervisor implementation; thus, they are considered restricted if an interception occurs.

When a nonconstrained transaction is aborted because of the attempted execution of a restricted instruction, the transaction abort code in the transaction diagnostic block is set to 11 (restricted instruction), and the condition code is set to 3, except as follows: when a nonconstrained transaction is aborted due to the attempted execution of an instruction that would otherwise result in a privileged operation exception, it is unpredictable whether the abort code is set to 11 (restricted instruction) or 4 (unfiltered program interruption resulting from the recognition of the privileged operation program interruption). When a nonconstrained transaction is aborted due to the attempted execution of PREFETCH DATA (RELATIVE LONG) when the code in the M_1 field is 6 or 7 or STORE CHARACTERS UNDER MASK HIGH when the M_3 field is zero and the code in the R_1 field is 6 or 7, it is unpredictable whether the abort code is set to 11 (restricted instruction) or 16 (cache other). When a nonconstrained transaction is aborted due to the attempted execution of MONITOR CALL, and both a monitor event condition and a specification exception condition are present it is unpredictable whether the abort code is set to 11 or 4, or, if the program interruption is filtered, 12.

Additional instructions may be restricted in a constrained transaction. Although these instructions are not currently defined to be restricted in a nonconstrained transaction, they may be restricted under certain circumstances in a nonconstrained transaction on future processors.

Certain restricted instructions may be allowed in the transactional execution mode on future processors. Therefore, the program should not rely on the transaction being aborted due to the attempted execution of a restricted instruction. The TRANSACTION ABORT instruction should be used to reliably cause a transaction to be aborted.

In a nonconstrained transaction, the program should provide an alternative non-transactional code path to accommodate a transaction that aborts due to a restricted instruction.

In operation, when the transaction nesting depth is zero, execution of the TRANSACTION BEGIN (TBEGIN) instruction resulting in condition code zero causes the CPU to enter the nonconstrained transactional execution mode. When the transaction nesting depth is zero, execution of the TRANSACTION BEGIN constrained (TBEGINC) instruction resulting in condition code zero causes the CPU to enter the constrained transactional execution mode.

Except where explicitly noted otherwise, all rules that apply for non-transactional execution also apply to transactional execution. Below are additional characteristics of processing while the CPU is in the transactional execution mode.

When the CPU is in the nonconstrained transactional execution mode, execution of the TRANSACTION BEGIN instruction resulting in condition code zero causes the CPU to remain in the nonconstrained transactional execution mode.

As observed by the CPU, fetches and stores made in the transaction execution mode are no different than those made while not in the transactional execution mode. As observed by other CPUs and by the I/O subsystem, all storage operand accesses made while a CPU is in the transactional execution mode appear to be a single block concurrent access. That is, the accesses to all bytes within a halfword, word, doubleword, or quadword are specified to appear to be block concurrent as observed by other CPUs and I/O (e.g., channel) programs. The halfword, word, doubleword, or quadword is referred to in this section as a block. When a fetch-type reference is specified to appear to be concurrent

within a block, no store access to the block by another CPU or I/O program is permitted during the time that bytes contained in the block are being fetched. When a store-type reference is specified to appear to be concurrent within a block, no access to the block, either fetch or store, is permitted by another CPU or I/O program during the time that the bytes within the block are being stored.

Storage accesses for instruction and DAT and ART (Access Register Table) table fetches follow the non-transactional rules.

The CPU leaves the transactional execution mode normally by means of a TRANSACTION END instruction that causes the transaction nesting depth to transition to zero or a CONDITIONAL TRANSACTION END instruction that sets the condition code to zero; in either of these cases, the transaction completes.

When the CPU leaves the transactional execution mode by means of the completion of a TRANSACTION END instruction or a CONDITIONAL TRANSACTION END instruction that sets the condition code to zero, all stores made while in the transactional execution mode are committed; that is, the stores appear to occur as a single block-concurrent operation as observed by other CPUs and by the I/O subsystem.

A transaction may be implicitly aborted for a variety of causes, or it may be explicitly aborted by the TRANSACTION ABORT instruction. Example possible causes of a transaction abort, the corresponding abort code, and the condition code that is placed into the transaction abort PSW are described below.

External Interruption: The transaction abort code is set to 2, and the condition code in the transaction abort PSW is set to 2. The transaction abort PSW is stored as the external old PSW as a part of external interruption processing.

Program Interruption (Unfiltered): A program interruption condition that results in an interruption (that is, an unfiltered condition) causes the transaction to be aborted with code 4. The condition code in the transaction abort PSW is set specific to the program interruption code. The transaction abort PSW is stored as the program old PSW as a part of program interruption processing.

An instruction that would otherwise result in a transaction being aborted due to an operation exception may yield alternate results: for a nonconstrained transaction, the transaction may instead abort with abort code 11 (restricted instruction); for a constrained transaction, a transaction constraint program interruption may be recognized instead of the operation exception.

When a PER (Program Event Recording) event is recognized in conjunction with any other unfiltered program exception condition, the condition code is set to 3.

Machine Check Interruption: The transaction abort code is set to 5, and the condition code in the transaction abort PSW is set to 2. The transaction abort PSW is stored as the machine check old PSW as a part of machine check interruption processing.

I/O Interruption: The transaction abort code is set to 6, and the condition code in the transaction abort PSW is set to 2. The transaction abort PSW is stored as the I/O old PSW as a part of I/O interruption processing.

Fetch Overflow: A fetch overflow condition is detected when the transaction attempts to fetch from more locations than the CPU supports. The transaction abort code is set to 7, and the condition code is set to either 2 or 3.

Store Overflow: A store overflow condition is detected when the transaction attempts to store to more locations than the CPU supports. The transaction abort code is set to 8, and the condition code is set to either 2 or 3.

Allowing the condition code to be either 2 or 3 in response to a fetch or store overflow abort allows the CPU to indicate potentially retryable situations (e.g., condition code 2 indicates re-execution of the transaction may be productive; while condition code 3 does not recommend re-execution).

Fetch Conflict: A fetch conflict condition is detected when another CPU or the I/O subsystem attempts to store into a location that has been transactionally fetched by this CPU. The transaction abort code is set to 9, and the condition code is set to 2.

Store Conflict: A store conflict condition is detected when another CPU or the I/O subsystem attempts to access a location that has been stored during transactional execution by this CPU. The transaction abort code is set to 10, and the condition code is set to 2.

Restricted Instruction: When the CPU is in the transactional execution mode, attempted execution of a restricted instruction causes the transaction to be aborted. The transaction abort code is set to 11, and the condition code is set to 3.

When the CPU is in the constrained transactional execution mode, it is unpredictable whether attempted execution of a restricted instruction results in a transaction constraint program interruption or an abort due to a restricted instruction. The transaction is still aborted but the abort code may indicate either cause.

Program Exception Condition (Filtered): A program exception condition that does not result in an interruption (that is, a filtered condition) causes the transaction to be aborted with a transaction abort code of 12. The condition code is set to 3.

Nesting Depth Exceeded: The nesting depth exceeded condition is detected when the transaction nesting depth is at the maximum allowable value for the configuration, and a TRANSACTION BEGIN instruction is executed. The transaction is aborted with a transaction abort code of 13, and the condition code is set to 3.

Cache Fetch Related Condition: A condition related to storage locations fetched by the transaction is detected by the CPU's cache circuitry. The transaction is aborted with a transaction abort code of 14, and the condition code is set to either 2 or 3.

Cache Store Related Condition: A condition related to storage locations stored by the transaction is detected by the CPU's cache circuitry. The transaction is aborted with a transaction abort code of 15, and the condition code is set to either 2 or 3.

Cache Other Condition: A cache other condition is detected by the CPU's cache circuitry. The transaction is aborted with a transaction abort code of 16, and the condition code is set to either 2 or 3.

During transactional execution, if the CPU accesses instructions or storage operands using different logical addresses that are mapped to the same absolute address, it is model dependent whether the transaction is aborted. If the transaction is aborted due to accesses using different logical addresses mapped to the same absolute address, abort code 14, 15, or 16 is set, depending on the condition.

CTEND Negative Operand Condition: A CTEND negative operand condition is recognized when the CONDITIONAL TRANSACTION END instruction is executed, and the contents of the second operand are negative. The

transaction is aborted with a transaction abort code of 17, and the condition code is set to 3.

CTEND Time Out Condition: A CTEND time out condition is recognized when the execution of a CONDITIONAL TRANSACTION END instruction exceeds a model-dependent limit. The transaction abort code is set to 18, and the condition code is set to 2.

Miscellaneous Condition: A miscellaneous condition is any other condition recognized by the CPU that causes the transaction to abort. The transaction abort code is set to 255 and the condition code is set to either 2 or 3.

When multiple configurations are executing in the same machine (for example, logical partitions or virtual machines), a transaction may be aborted due to an external machine check or I/O interruption that occurred in a different configuration.

Although examples are provided above, other causes of a transaction abort with corresponding abort codes and condition codes may be provided. For instance, a cause may be a Restart Interruption, in which the transaction abort code is set to 1, and the condition code in the transaction abort PSW is set to 2. The transaction abort PSW is stored as the restart-old PSW as a part of restart processing. As a further example, a cause may be a Supervisor Call condition, in which the abort code is set to 3, and the condition code in the transaction abort PSW is set to 3. Other or different examples are also possible.

Notes:

1. The miscellaneous condition may result from any of the following:

Instructions, such as, in the z/Architecture, COMPARE AND REPLACE DAT TABLE ENTRY, COMPARE AND SWAP AND PURGE, INVALIDATE DAT TABLE ENTRY, INVALIDATE PAGE TABLE ENTRY, PERFORM FRAME MANAGEMENT FUNCTION in which the NQ control is zero and the SK control is one, SET STORAGE KEY EXTENDED in which the NQ control is zero, performed by another CPU in the configuration; the condition code is set to 2.

An operator function, such as reset, restart or stop, or the equivalent SIGNAL PROCESSOR order is performed on the CPU. The condition code is set to 2. Any other condition not enumerated above; the condition code is set to 2 or 3.

2. The location at which fetch and store conflicts are detected may be anywhere within the same cache line.
3. Under certain conditions, the CPU may not be able to distinguish between similar abort conditions. For example, a fetch or store overflow may be indistinguishable from a respective fetch or store conflict.
4. Speculative execution of multiple instruction paths by the CPU may result in a transaction being aborted due to conflict or overflow conditions, even if such conditions do not occur in the conceptual sequence. While in the constrained transactional execution mode, the CPU may temporarily inhibit speculative execution, allowing the transaction to attempt to complete without detecting such conflicts or overflows speculatively.

Execution of a TRANSACTION ABORT instruction causes the transaction to abort. The transaction abort code is set from the second operand address. The condition code is set to either 2 or 3, depending on whether bit 63 of the second operand address is zero or one, respectively.

FIG. 8 depicts a workflow 800 when program code executes a transaction in a computer system. This workflow 800 provides an additional illustration to the discussion of

transactions provided above. As illustrated in FIG. 8, program code begins a transaction and this transaction begins with a TRANSACTION BEGIN instruction (810). For example, a constrained transaction begins with a TBEGINC instructions and, as with a nonconstrained transaction, ends normally with a TEND instruction, though, as discussed, the transaction may abort for many system-defined reasons. The program code sets a transaction indicator, such as a bit (820) and fetches multiple sequential instructions (830). The program code decodes the instructions (840), issues the instructions (850), and executes the instructions. If the program code has reached the instruction specifying the transaction end (860), in one example, a TRANSACTION END (TEND) instruction, the program code saves the result (870) and resets the transaction indicator (e.g., bit) (880). If there is no TRANSACTION END (TEND) instruction, the program code returns to fetching multiple sequential instructions (830).

FIG. 9 depicts an example of how program code may abort a transaction when there is a conflict. The issue illustrated in FIG. 9 is mitigated by aspects of some embodiments of the present invention. The transaction depicted in FIG. 9 is a constrained transaction, as indicated by the processor's attempt to execute the transactions after aborting the transaction based on a conflict. Referring to FIG. 9, the program code initiates a transaction (910). If the program code detects a conflict (920), the program code aborts the transaction (930a) and will attempt to execute the transaction (910) at a different time. If the program code does not detect a conflict (920), the program code will execute the instructions in the transaction (930b) and save the result (940b).

FIG. 10 illustrates certain aspects of the functionality of a system where the super core indicator of certain embodiments of the present invention is utilized in order to ensure the completion of a given transaction. As illustrated in FIG. 10, the program code initiates a transaction (1010). If the program code does not detect a conflict (1020), the program code executes the transaction (i.e., the instructions in the transaction) (1030) and saves results the results (1040). However, if the program code detects a conflict (1020), the program code checks for a super core indicator (1025). As aforementioned, each core may include an indicator, for example, a bit, including but not limited to, a super core bit that the program code utilizes to designate a core, a super core, and/or to determine whether a core is a super core. If the super core indicator indicates that the processor executing the transaction is not a super core, the program code aborts the transaction (1027), and the program code will attempt to initiate the transaction again (1010). If the super core indicator indicates that the processor executing the transaction is a super core, the program code executes the transaction (1030), saves results the results (1040), and clears the super core indicator (1050).

When a core is designated a super core, the program code will complete execution of a constrained transaction (1030) because even if there is a conflict, for example, during execution, the super core receives a request for data it contains in its cache which are marked as a read or write data, the program code will not abort the transaction. In an embodiment of the present invention, in response to receiving a request for data, a designated super core will ignore the request for data, it will not provide access to the data, but will instead continue to execute the transaction to its completion. In an embodiment of the present invention, if the program code obtains a request for data utilized by the super core executing the transaction, it will not respond to

the request. In another embodiment of the present invention, the program code will respond to the request with a pre-defined "rejecting response." In another embodiment of the present invention, the program code queues pending request and once the transaction completes, the program code respond to and/or fulfills the requests.

In an embodiment of the present invention, the program code clears the super core indicator (1050) following the execution of a transaction (1030) so that in the system, only a single super core is active at a time. By ending the super core status with the completion of a transaction, the program code ensures a "fairness" among, for example, multiple cores competing for this preferential transactional status. In an embodiment of the present invention, the program code enables a single super core within each partition.

To avoid multiple super cores within a given environment, an embodiment of the present invention includes an arbitration cycle, which the program code utilizes to turn one core into a super core, to avoid multiple simultaneous requests turning multiple cores concurrently into super cores. This arbitration aspect, which is included in some embodiments of the present invention, is illustrated by FIG. 11.

Referring to FIG. 11, in an embodiment of the present invention, program code obtains a request for super core status from an arbitrator program code in a system (1110). Based on this request, the program code checks for a super core indicator in the system (1120). If there is already a super core in the system, a super core indicator is set, and the arbitrator program code refuses the request (1130a). However, if there is no super core set in the system, which is identifiable with the indicator indicating this status, the arbitrator program code grants the status of a super core to the core which requested this status, and sets the indicator to the arbitrator core (1130b). The program code then monitors the core as it executes a transaction (1140b). Based on the super core committing the transaction, the arbitrator program code clears the super core (1150b). By clearing this indicator, should the arbitrator program code receive a subsequent request for super core status, it will be able to grant the status as no core in the system will have the super core indicator.

As illustrated in FIG. 12, in an embodiment of the present invention, program code may set a super core value that is associated with a transaction to ensure completion of the transaction after that transaction has already failed once or multiple times. In this manner, the program code avoids attempting to re-execute and aborting the same transaction repeatedly. In FIG. 12, the program code initiates a transaction (1210), determines there is a conflict (1215) and based on determining that there is a conflict, determines whether there is a super core indicator associated with the core executing the transaction (1220). If there is a super core indicator, as discussed above, the program code completes the transaction (1230a) and clears the super core indicator (1240a). If there is no super core indicator associated with the core executing this transaction, the program code aborts the transaction (1230b) and based on aborting the transaction, the program code modifies a super core value associated with the transaction (1240b). In an embodiment of the present invention, the program code modifies the super core value associates with the transaction by setting the super core value. In another embodiment of the present invention, the program code modifies the super core value by incrementing the super core value for each failed transaction. In some embodiments of the present invention, the super core value has to achieve pre-determined threshold in order to qualify to become a super core. In an embodiment of the

present invention, this value is one. In another embodiment, this value is a number larger than one. In one embodiment, the super core value is set by a program code. In another embodiment, this value is hard-coded setting in a hardware and/or a software element of a computer system. Thus, when the core attempts again to execute this transaction, the transaction is ensured to complete. The setting of the super core value is further explained in FIG. 13.

In an embodiment of the present invention, each time a transaction fails, the program code may increment a super core value associated with the transaction by one. As seen in FIG. 13, program code initiates a transaction and as part of initiating the transaction, obtains a super core value associated with a transaction to determine whether the super core value is greater than or equal to a threshold value (1310). In an embodiment of the present invention, this threshold value may be predefined, hard-coded or configurable, static or dynamic. Based on determining that the super core value is greater than or equal to the threshold, the program code requests super core status for the core executing the transaction (1320). The program code obtains the super core status (once available) and the program code sets a super core indicator in the core (1330).

FIG. 14 is a workflow 1400 that depicts some aspects of some embodiments of the present invention. In an embodiment of the present invention, a processor initiates an execution of a transaction, where the processor comprises a transactional core, and where the execution of the transaction is performed by the transactional core (1410). In an embodiment of the present invention, the transaction is a constrained transaction (1412).

The processor obtains, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional core in the computing environment (1420). In an embodiment of the present invention, the indication comprises a request from the additional core related to data utilized by the transactional core during the execution of the transaction (1422). In this embodiment, the data may be located in a cache of the transactional core and is marked as read data or write data (1424). In an embodiment of the present invention, the processor queues the request and based on completing executing the transaction, the processor fulfills the request, by providing access to the data to the additional core (1426).

In an embodiment of the present invention, the processor determines if the transactional core comprises an indicator (1430).

Based on determining that the transactional core comprises an indicator, the processor ignores the conflict and utilizes the transactional core to complete executing the transaction (1440a). In an embodiment of the present invention where the indication comprises a request from the additional core related to data utilized by the transactional core during the execution of the transaction the ignoring may include the processor denying the additional core access to the data (1442).

In an embodiment of the present invention, the processor monitors the executing of the transaction by the transactional core (1450) and based on the transactional core completing the executing by committing the transaction, the processor clears the indicator (1460).

In an embodiment of the present invention, the processor obtains a request from the transactional core for the indicator (1402). The processor determines if a super core status is available (1404). Based on determining that a super core

status is available, the processor modifies the transactional core, to provide the indicator (1406).

In an embodiment of the present invention, based on determining that the transactional core does not comprise an indicator, the processor aborts the transaction (1441). The processor increments a value associated with the transaction (1443).

In an embodiment of the present invention, the processor obtains the value associated with the transaction (1445) and the processor evaluates the value to determine if the value meets a threshold (1447). Based on determining that the value meets the threshold, the processor requests an indicator for the transactional core (1449). The processor determines if a super core status is available (1404). Based on determining that a super core status is available, the processor modifies the transactional core, to provide the indicator (1406).

In an embodiment of the present invention, the processor determines that the super core status is available by accessing cores in the computing environment to determine if a threshold number of cores comprise indicators (1405); and determining that the super core status is available based on less than the threshold number of cores comprising indicators (1407).

The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

Referring to FIG. 15, in one example, a computer program product 1200 includes, for instance, one or more non-transitory computer readable storage media 1202 to store computer readable program code means, logic and/or instructions 1204 thereon to provide and facilitate one or more embodiments.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may com-

prise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on

the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

In addition to the above, one or more aspects may be provided, offered, deployed, managed, serviced, etc. by a service provider who offers management of customer environments. For instance, the service provider can create, maintain, support, etc. computer code and/or a computer infrastructure that performs one or more aspects for one or more customers. In return, the service provider may receive payment from the customer under a subscription and/or fee agreement, as examples. Additionally or alternatively, the service provider may receive payment from the sale of advertising content to one or more third parties.

In one aspect, an application may be deployed for performing one or more embodiments. As one example, the deploying of an application comprises providing computer infrastructure operable to perform one or more embodiments.

As a further aspect, a computing infrastructure may be deployed comprising integrating computer readable code into a computing system, in which the code in combination with the computing system is capable of performing one or more embodiments.

As yet a further aspect, a process for integrating computing infrastructure comprising integrating computer readable code into a computer system may be provided. The computer system comprises a computer readable medium, in which the computer medium comprises one or more embodiments. The code in combination with the computer system is capable of performing one or more embodiments.

Although various embodiments are described above, these are only examples. For example, computing environments of other architectures can be used to incorporate and use one or more embodiments. Further, different instructions, instruction formats, instruction fields and/or instruction values may be used. Yet further, although examples of values for abort codes and condition codes are provided, other values may be used. Moreover, different, other, and/or additional restrictions/constraints may be provided/used. Yet further, other intervals may be provided and/or used in differing ways. Many variations are possible.

Further, other types of computing environments can benefit and be used. As an example, a data processing system suitable for storing and/or executing program code is usable that includes at least two processors coupled directly or indirectly to memory elements through a system bus. The

memory elements include, for instance, local memory employed during actual execution of the program code, bulk storage, and cache memory which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

Input/Output or I/O devices (including, but not limited to, keyboards, displays, pointing devices, DASD, tape, CDs, DVDs, thumb drives and other memory media, etc.) can be coupled to the system either directly or through intervening I/O controllers. Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modems, and Ethernet cards are just a few of the available types of network adapters.

Referring to FIG. 16, representative components of a Host Computer system **5000** to implement one or more embodiments are portrayed. The representative host computer **5000** comprises one or more CPUs **5001** in communication with computer memory (i.e., central storage) **5002**, as well as I/O interfaces to storage media devices **5011** and networks **5010** for communicating with other computers or SANs and the like. The CPU **5001** is compliant with an architecture having an architected instruction set and architected functionality. The CPU **5001** may have access register translation (ART) **5012**, which includes an ART lookaside buffer (ALB) **5013**, for selecting an address space to be used by dynamic address translation (DAT) **5003** for transforming program addresses (virtual addresses) into real addresses of memory. A DAT typically includes a translation lookaside buffer (TLB) **5007** for caching translations so that later accesses to the block of computer memory **5002** do not require the delay of address translation. Typically, a cache **5009** is employed between computer memory **5002** and the processor **5001**. The cache **5009** may be hierarchical having a large cache available to more than one CPU and smaller, faster (lower level) caches between the large cache and each CPU. In some implementations, the lower level caches are split to provide separate low level caches for instruction fetching and data accesses. In one embodiment, for the TX facility, a transaction diagnostic block (TDB) **5100** and one or more buffers **5101** may be stored in one or more of cache **5009** and memory **5002**. In one example, in TX mode, data is initially stored in a TX buffer, and when TX mode ends (e.g., outermost TEND), the data in the buffer is stored (committed) to memory, or if there is an abort, the data in the buffer is discarded.

In one embodiment, an instruction is fetched from memory **5002** by an instruction fetch unit **5004** via a cache **5009**. The instruction is decoded in an instruction decode unit **5006** and dispatched (with other instructions in some embodiments) to instruction execution unit or units **5008**. Typically several execution units **5008** are employed, for example an arithmetic execution unit, a floating point execution unit and a branch instruction execution unit. Further, in one embodiment of the TX facility, various TX controls **5110** may be employed. The instruction is executed by the execution unit, accessing operands from instruction specified registers or memory as needed. If an operand is to be accessed (loaded or stored) from memory **5002**, a load/store unit **5005** typically handles the access under control of the instruction being executed. Instructions may be executed in hardware circuits or in internal microcode (firmware) or by a combination of both.

In accordance with an aspect of the TX facility, processor **5001** also includes a PSW **5102** (e.g., TX and/or abort PSW), a nesting depth **5104**, a TDBA **5106**, and one or more control registers **5108**.

As noted, a computer system includes information in local (or main) storage, as well as addressing, protection, and reference and change recording. Some aspects of addressing include the format of addresses, the concept of address spaces, the various types of addresses, and the manner in which one type of address is translated to another type of address. Some of main storage includes permanently assigned storage locations. Main storage provides the system with directly addressable fast-access storage of data. Both data and programs are to be loaded into main storage (from input devices) before they can be processed.

Main storage may include one or more smaller, faster-access buffer storages, sometimes called caches. A cache is typically physically associated with a CPU or an I/O processor. The effects, except on performance, of the physical construction and use of distinct storage media are generally not observable by the program.

Separate caches may be maintained for instructions and for data operands. Information within a cache is maintained in contiguous bytes on an integral boundary called a cache block or cache line (or line, for short). A model may provide an EXTRACT CACHE ATTRIBUTE instruction which returns the size of a cache line in bytes. A model may also provide PREFETCH DATA and PREFETCH DATA RELATIVE LONG instructions which effects the prefetching of storage into the data or instruction cache or the releasing of data from the cache.

Storage is viewed as a long horizontal string of bits. For most operations, accesses to storage proceed in a left-to-right sequence. The string of bits is subdivided into units of eight bits. An eight-bit unit is called a byte, which is the basic building block of all information formats. Each byte location in storage is identified by a unique nonnegative integer, which is the address of that byte location or, simply, the byte address. Adjacent byte locations have consecutive addresses, starting with 0 on the left and proceeding in a left-to-right sequence. Addresses are unsigned binary integers and are 24, 31, or 64 bits.

Information is transmitted between storage and a CPU or a channel subsystem one byte, or a group of bytes, at a time. Unless otherwise specified, in, for instance, the z/Architecture, a group of bytes in storage is addressed by the leftmost byte of the group. The number of bytes in the group is either implied or explicitly specified by the operation to be performed. When used in a CPU operation, a group of bytes is called a field. Within each group of bytes, in, for instance, the z/Architecture, bits are numbered in a left-to-right sequence. In the z/Architecture, the leftmost bits are sometimes referred to as the "high-order" bits and the rightmost bits as the "low-order" bits. Bit numbers are not storage addresses, however. Only bytes can be addressed. To operate on individual bits of a byte in storage, the entire byte is accessed. The bits in a byte are numbered 0 through 7, from left to right (in, e.g., the z/Architecture). The bits in an address may be numbered 8-31 or 40-63 for 24-bit addresses, or 1-31 or 33-63 for 31-bit addresses; they are numbered 0-63 for 64-bit addresses. In one example, bits 8-31 and 1-31 apply to addresses that are in a location (e.g., register) that is 32 bits wide, whereas bits 40-63 and 33-63 apply to addresses that are in a 64-bit wide location. Within any other fixed-length format of multiple bytes, the bits making up the format are consecutively numbered starting from 0. For purposes of error detection, and in preferably for

correction, one or more check bits may be transmitted with each byte or with a group of bytes. Such check bits are generated automatically by the machine and cannot be directly controlled by the program. Storage capacities are expressed in number of bytes. When the length of a storage-operand field is implied by the operation code of an instruction, the field is said to have a fixed length, which can be one, two, four, eight, or sixteen bytes. Larger fields may be implied for some instructions. When the length of a storage-operand field is not implied but is stated explicitly, the field is said to have a variable length. Variable-length operands can vary in length by increments of one byte (or with some instructions, in multiples of two bytes or other multiples). When information is placed in storage, the contents of only those byte locations are replaced that are included in the designated field, even though the width of the physical path to storage may be greater than the length of the field being stored.

Certain units of information are to be on an integral boundary in storage. A boundary is called integral for a unit of information when its storage address is a multiple of the length of the unit in bytes. Special names are given to fields of 2, 4, 8, 16, and 32 bytes on an integral boundary. A halfword is a group of two consecutive bytes on a two-byte boundary and is the basic building block of instructions. A word is a group of four consecutive bytes on a four-byte boundary. A doubleword is a group of eight consecutive bytes on an eight-byte boundary. A quadword is a group of 16 consecutive bytes on a 16-byte boundary. An octoword is a group of 32 consecutive bytes on a 32-byte boundary. When storage addresses designate halfwords, words, doublewords, quadwords, and octowords, the binary representation of the address contains one, two, three, four, or five rightmost zero bits, respectively. Instructions are to be on two-byte integral boundaries. The storage operands of most instructions do not have boundary-alignment requirements.

On devices that implement separate caches for instructions and data operands, a significant delay may be experienced if the program stores into a cache line from which instructions are subsequently fetched, regardless of whether the store alters the instructions that are subsequently fetched.

In one example, the embodiment may be practiced by software (sometimes referred to licensed internal code, firmware, micro-code, milli-code, pico-code and the like, any of which would be consistent with one or more embodiments). Referring to FIG. 15, software program code which embodies one or more aspects may be accessed by processor 5001 of the host system 5000 from long-term storage media devices 5011, such as a CD-ROM drive, tape drive or hard drive. The software program code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from computer memory 5002 or storage of one computer system over a network 5010 to other computer systems for use by users of such other systems.

The software program code includes an operating system which controls the function and interaction of the various computer components and one or more application programs. Program code is normally paged from storage media device 5011 to the relatively higher-speed computer storage 5002 where it is available for processing by processor 5001. The techniques and methods for embodying software program code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein. Program code, when created and stored on a tangible medium (including but not limited to

electronic memory modules (RAM), flash memory, Compact Discs (CDs), DVDs, Magnetic Tape and the like is often referred to as a "computer program product". The computer program product medium is typically readable by a processing circuit preferably in a computer system for execution by the processing circuit.

FIG. 17 illustrates a representative workstation or server hardware system in which one or more embodiments may be practiced. The system 5020 of FIG. 17 comprises a representative base computer system 5021, such as a personal computer, a workstation or a server, including optional peripheral devices. The base computer system 5021 includes one or more processors 5026 and a bus employed to connect and enable communication between the processor(s) 5026 and the other components of the system 5021 in accordance with known techniques. The bus connects the processor 5026 to memory 5025 and long-term storage 5027 which can include a hard drive (including any of magnetic media, CD, DVD and Flash Memory for example) or a tape drive for example. The system 5021 might also include a user interface adapter, which connects the microprocessor 5026 via the bus to one or more interface devices, such as a keyboard 5024, a mouse 5023, a printer/scanner 5030 and/or other interface devices, which can be any user interface device, such as a touch sensitive screen, digitized entry pad, etc. The bus also connects a display device 5022, such as an LCD screen or monitor, to the microprocessor 5026 via a display adapter.

The system 5021 may communicate with other computers or networks of computers by way of a network adapter capable of communicating 5028 with a network 5029. Example network adapters are communications channels, token ring, Ethernet or modems. Alternatively, the system 5021 may communicate using a wireless interface, such as a CDPD (cellular digital packet data) card. The system 5021 may be associated with such other computers in a Local Area Network (LAN) or a Wide Area Network (WAN), or the system 5021 can be a client in a client/server arrangement with another computer, etc. All of these configurations, as well as the appropriate communications hardware and software, are known in the art.

FIG. 18 illustrates a data processing network 5040 in which one or more embodiments may be practiced. The data processing network 5040 may include a plurality of individual networks, such as a wireless network and a wired network, each of which may include a plurality of individual workstations 5041, 5042, 5043, 5044. Additionally, as those skilled in the art will appreciate, one or more LANs may be included, where a LAN may comprise a plurality of intelligent workstations coupled to a host processor.

Still referring to FIG. 18, the networks may also include mainframe computers or servers, such as a gateway computer (client server 5046) or application server (remote server 5048 which may access a data repository and may also be accessed directly from a workstation 5045). A gateway computer 5046 serves as a point of entry into each individual network. A gateway is needed when connecting one networking protocol to another. The gateway 5046 may be preferably coupled to another network (the Internet 5047 for example) by means of a communications link. The gateway 5046 may also be directly coupled to one or more workstations 5041, 5042, 5043, 5044 using a communications link. The gateway computer may be implemented utilizing an IBM eServer System z server available from International Business Machines Corporation.

Referring concurrently to FIG. 17 and FIG. 18, software programming code 5031 which may embody one or more

aspects may be accessed by the processor **5026** of the system **5020** from long-term storage media **5027**, such as a CD-ROM drive or hard drive. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users **5050**, **5051** from the memory or storage of one computer system over a network to other computer systems for use by users of such other systems.

Alternatively, the programming code may be embodied in the memory **5025**, and accessed by the processor **5026** using the processor bus. Such programming code includes an operating system which controls the function and interaction of the various computer components and one or more application programs **5032**. Program code is normally paged from storage media **5027** to high-speed memory **5025** where it is available for processing by the processor **5026**. The techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks are well known and will not be further discussed herein. Program code, when created and stored on a tangible medium (including but not limited to electronic memory modules (RAM), flash memory, Compact Discs (CDs), DVDs, Magnetic Tape and the like is often referred to as a "computer program product". The computer program product medium is typically readable by a processing circuit preferably in a computer system for execution by the processing circuit.

The cache that is most readily available to the processor (normally faster and smaller than other caches of the processor) is the lowest (L1 or level one) cache and main store (main memory) is the highest level cache (L3 if there are 3 levels). The lowest level cache is often divided into an instruction cache (I-Cache) holding machine instructions to be executed and a data cache (D-Cache) holding data operands.

Referring to FIG. 19, an exemplary processor embodiment is depicted for processor **5026**. Typically one or more levels of cache **5053** are employed to buffer memory blocks in order to improve processor performance. The cache **5053** is a high speed buffer holding cache lines of memory data that are likely to be used. Typical cache lines are 64, 128 or 256 bytes of memory data. Separate caches are often employed for caching instructions than for caching data. Cache coherence (synchronization of copies of lines in memory and the caches) is often provided by various "snoop" algorithms well known in the art. Main memory storage **5025** of a processor system is often referred to as a cache. In a processor system having 4 levels of cache **5053**, main storage **5025** is sometimes referred to as the level 5 (L5) cache since it is typically faster and only holds a portion of the non-volatile storage (DASD, tape etc) that is available to a computer system. Main storage **5025** "caches" pages of data paged in and out of the main storage **5025** by the operating system.

A program counter (instruction counter) **5061** keeps track of the address of the current instruction to be executed. A program counter in a z/Architecture processor is 64 bits and can be truncated to 31 or 24 bits to support prior addressing limits. A program counter is typically embodied in a PSW (program status word) of a computer such that it persists during context switching. Thus, a program in progress, having a program counter value, may be interrupted by, for example, the operating system (context switch from the program environment to the operating system environment). The PSW of the program maintains the program counter value while the program is not active, and the program

counter (in the PSW) of the operating system is used while the operating system is executing. Typically, the program counter is incremented by an amount equal to the number of bytes of the current instruction. RISC (Reduced Instruction Set Computing) instructions are typically fixed length while CISC (Complex Instruction Set Computing) instructions are typically variable length. Instructions of the IBM z/Architecture are CISC instructions having a length of 2, 4 or 6 bytes. The Program counter **5061** is modified by either a context switch operation or a branch taken operation of a branch instruction for example. In a context switch operation, the current program counter value is saved in the program status word along with other state information about the program being executed (such as condition codes), and a new program counter value is loaded pointing to an instruction of a new program module to be executed. A branch taken operation is performed in order to permit the program to make decisions or loop within the program by loading the result of the branch instruction into the program counter **5061**.

Typically an instruction fetch unit **5055** is employed to fetch instructions on behalf of the processor **5026**. The fetch unit either fetches "next sequential instructions", target instructions of branch taken instructions, or first instructions of a program following a context switch. Modern Instruction fetch units often employ prefetch techniques to speculatively prefetch instructions based on the likelihood that the prefetched instructions might be used. For example, a fetch unit may fetch 16 bytes of instruction that includes the next sequential instruction and additional bytes of further sequential instructions.

The fetched instructions are then executed by the processor **5026**. In an embodiment, the fetched instruction(s) are passed to a dispatch unit **5056** of the fetch unit. The dispatch unit decodes the instruction(s) and forwards information about the decoded instruction(s) to appropriate units **5057**, **5058**, **5060**. An execution unit **5057** will typically receive information about decoded arithmetic instructions from the instruction fetch unit **5055** and will perform arithmetic operations on operands according to the opcode of the instruction. Operands are provided to the execution unit **5057** preferably either from memory **5025**, architected registers **5059** or from an immediate field of the instruction being executed. Results of the execution, when stored, are stored either in memory **5025**, registers **5059** or in other machine hardware (such as control registers, PSW registers and the like).

Virtual addresses are transformed into real addresses using dynamic address translation **5062** and, optionally, using access register translation **5063**.

A processor **5026** typically has one or more units **5057**, **5058**, **5060** for executing the function of the instruction. Referring to FIG. 20A, an execution unit **5057** may communicate **5071** with architected general registers **5059**, a decode/dispatch unit **5056**, a load store unit **5060**, and other **5065** processor units by way of interfacing logic **5071**. An execution unit **5057** may employ several register circuits **5067**, **5068**, **5069** to hold information that the arithmetic logic unit (ALU) **5066** will operate on. The ALU performs arithmetic operations such as add, subtract, multiply and divide as well as logical function such as and, or and exclusive-or (XOR), rotate and shift. Preferably the ALU supports specialized operations that are design dependent. Other circuits may provide other architected facilities **5072** including condition codes and recovery support logic for example. Typically the result of an ALU operation is held in an output register circuit **5070** which can forward the result

to a variety of other processing functions. There are many arrangements of processor units, the present description is only intended to provide a representative understanding of one embodiment.

An ADD instruction for example would be executed in an execution unit **5057** having arithmetic and logical functionality while a floating point instruction for example would be executed in a floating point execution having specialized floating point capability. Preferably, an execution unit operates on operands identified by an instruction by performing an opcode defined function on the operands. For example, an ADD instruction may be executed by an execution unit **5057** on operands found in two registers **5059** identified by register fields of the instruction.

The execution unit **5057** performs the arithmetic addition on two operands and stores the result in a third operand where the third operand may be a third register or one of the two source registers. The execution unit preferably utilizes an Arithmetic Logic Unit (ALU) **5066** that is capable of performing a variety of logical functions such as Shift, Rotate, And, Or and XOR as well as a variety of algebraic functions including any of add, subtract, multiply, divide. Some ALUs **5066** are designed for scalar operations and some for floating point. Data may be Big Endian (where the least significant byte is at the highest byte address) or Little Endian (where the least significant byte is at the lowest byte address) depending on architecture. The IBM z/Architecture is Big Endian. Signed fields may be sign and magnitude, 1's complement or 2's complement depending on architecture. A 2's complement number is advantageous in that the ALU does not need to design a subtract capability since either a negative value or a positive value in 2's complement requires only an addition within the ALU. Numbers are commonly described in shorthand, where a 12 bit field defines an address of a 4,096 byte block and is commonly described as a 4 Kbyte (Kilo-byte) block, for example.

Referring to FIG. 20B, branch instruction information for executing a branch instruction is typically sent to a branch unit **5058** which often employs a branch prediction algorithm such as a branch history table **5082** to predict the outcome of the branch before other conditional operations are complete. The target of the current branch instruction will be fetched and speculatively executed before the conditional operations are complete. When the conditional operations are completed the speculatively executed branch instructions are either completed or discarded based on the conditions of the conditional operation and the speculated outcome. A typical branch instruction may test condition codes and branch to a target address if the condition codes meet the branch requirement of the branch instruction, a target address may be calculated based on several numbers including ones found in register fields or an immediate field of the instruction for example. The branch unit **5058** may employ an ALU **5074** having a plurality of input register circuits **5075**, **5076**, **5077** and an output register circuit **5080**. The branch unit **5058** may communicate **5081** with general registers **5059**, decode dispatch unit **5056** or other circuits **5073**, for example.

The execution of a group of instructions can be interrupted for a variety of reasons including a context switch initiated by an operating system, a program exception or error causing a context switch, an I/O interruption signal causing a context switch or multi-threading activity of a plurality of programs (in a multi-threaded environment), for example. Preferably a context switch action saves state information about a currently executing program and then loads state information about another program being

invoked. State information may be saved in hardware registers or in memory for example. State information preferably comprises a program counter value pointing to a next instruction to be executed, condition codes, memory translation information and architected register content. A context switch activity can be exercised by hardware circuits, application programs, operating system programs or firmware code (microcode, pico-code or licensed internal code (LIC)) alone or in combination.

A processor accesses operands according to instruction defined methods. The instruction may provide an immediate operand using the value of a portion of the instruction, may provide one or more register fields explicitly pointing to either general purpose registers or special purpose registers (floating point registers for example). The instruction may utilize implied registers identified by an opcode field as operands. The instruction may utilize memory locations for operands. A memory location of an operand may be provided by a register, an immediate field, or a combination of registers and immediate field as exemplified by the z/Architecture long displacement facility wherein the instruction defines a base register, an index register and an immediate field (displacement field) that are added together to provide the address of the operand in memory for example. Location herein typically implies a location in main memory (main storage) unless otherwise indicated.

Referring to FIG. 20C, a processor accesses storage using a load/store unit **5060**. The load/store unit **5060** may perform a load operation by obtaining the address of the target operand in memory **5053** and loading the operand in a register **5059** or another memory **5053** location, or may perform a store operation by obtaining the address of the target operand in memory **5053** and storing data obtained from a register **5059** or another memory **5053** location in the target operand location in memory **5053**. The load/store unit **5060** may be speculative and may access memory in a sequence that is out-of-order relative to instruction sequence, however the load/store unit **5060** is to maintain the appearance to programs that instructions were executed in order. A load/store unit **5060** may communicate **5084** with general registers **5059**, decode/dispatch unit **5056**, cache/memory interface **5053** or other elements **5083** and comprises various register circuits **5086**, **5087**, **5088** and **5089**, ALUs **5085** and control logic **5090** to calculate storage addresses and to provide pipeline sequencing to keep operations in-order. Some operations may be out of order but the load/store unit provides functionality to make the out of order operations to appear to the program as having been performed in order, as is well known in the art.

Preferably addresses that an application program "sees" are often referred to as virtual addresses. Virtual addresses are sometimes referred to as "logical addresses" and "effective addresses". These virtual addresses are virtual in that they are redirected to physical memory location by one of a variety of dynamic address translation (DAT) technologies including, but not limited to, simply prefixing a virtual address with an offset value, translating the virtual address via one or more translation tables, the translation tables preferably comprising at least a segment table and a page table alone or in combination, preferably, the segment table having an entry pointing to the page table. In the z/Architecture, a hierarchy of translation is provided including a region first table, a region second table, a region third table, a segment table and an optional page table. The performance of the address translation is often improved by utilizing a translation lookaside buffer (TLB) which comprises entries mapping a virtual address to an associated physical memory

location. The entries are created when the DAT translates a virtual address using the translation tables. Subsequent use of the virtual address can then utilize the entry of the fast TLB rather than the slow sequential translation table accesses. TLB content may be managed by a variety of replacement algorithms including LRU (Least Recently used).

In the case where the processor is a processor of a multi-processor system, each processor has responsibility to keep shared resources, such as I/O, caches, TLBs and memory, interlocked for coherency. Typically, "snoop" technologies will be utilized in maintaining cache coherency. In a snoop environment, each cache line may be marked as being in any one of a shared state, an exclusive state, a changed state, an invalid state and the like in order to facilitate sharing.

I/O units **5054** (FIG. 19) provide the processor with means for attaching to peripheral devices including tape, disc, printers, displays, and networks for example. I/O units are often presented to the computer program by software drivers. In mainframes, such as the System z from IBM®, channel adapters and open system adapters are I/O units of the mainframe that provide the communications between the operating system and peripheral devices.

Further, other types of computing environments can benefit from one or more aspects. As an example, an environment may include an emulator (e.g., software or other emulation mechanisms), in which a particular architecture (including, for instance, instruction execution, architected functions, such as address translation, and architected registers) or a subset thereof is emulated (e.g., on a native computer system having a processor and memory). In such an environment, one or more emulation functions of the emulator can implement one or more embodiments, even though a computer executing the emulator may have a different architecture than the capabilities being emulated. As one example, in emulation mode, the specific instruction or operation being emulated is decoded, and an appropriate emulation function is built to implement the individual instruction or operation.

In an emulation environment, a host computer includes, for instance, a memory to store instructions and data; an instruction fetch unit to fetch instructions from memory and to optionally, provide local buffering for the fetched instruction; an instruction decode unit to receive the fetched instructions and to determine the type of instructions that have been fetched; and an instruction execution unit to execute the instructions. Execution may include loading data into a register from memory; storing data back to memory from a register; or performing some type of arithmetic or logical operation, as determined by the decode unit. In one example, each unit is implemented in software. For instance, the operations being performed by the units are implemented as one or more subroutines within emulator software.

More particularly, in a mainframe, architected machine instructions are used by programmers, usually today "C" programmers, often by way of a compiler application. These instructions stored in the storage medium may be executed natively in a z/Architecture IBM® Server, or alternatively in machines executing other architectures. They can be emulated in the existing and in future IBM mainframe servers and on other machines of IBM® (e.g., Power Systems servers and System x Servers). They can be executed in machines running Linux on a wide variety of machines using hardware manufactured by IBM®, Intel®, AMD, and others. Besides execution on that hardware under a z/Architecture, Linux can be used as well as machines which use

emulation by Hercules, UMX, or FSI (Fundamental Software, Inc), where generally execution is in an emulation mode. In emulation mode, emulation software is executed by a native processor to emulate the architecture of an emulated processor.

The native processor typically executes emulation software comprising either firmware or a native operating system to perform emulation of the emulated processor. The emulation software is responsible for fetching and executing instructions of the emulated processor architecture. The emulation software maintains an emulated program counter to keep track of instruction boundaries. The emulation software may fetch one or more emulated machine instructions at a time and convert the one or more emulated machine instructions to a corresponding group of native machine instructions for execution by the native processor. These converted instructions may be cached such that a faster conversion can be accomplished. Notwithstanding, the emulation software is to maintain the architecture rules of the emulated processor architecture so as to assure operating systems and applications written for the emulated processor operate correctly. Furthermore, the emulation software is to provide resources identified by the emulated processor architecture including, but not limited to, control registers, general purpose registers, floating point registers, dynamic address translation function including segment tables and page tables for example, interrupt mechanisms, context switch mechanisms, Time of Day (TOD) clocks and architected interfaces to I/O subsystems such that an operating system or an application program designed to run on the emulated processor, can be run on the native processor having the emulation software.

A specific instruction being emulated is decoded, and a subroutine is called to perform the function of the individual instruction. An emulation software function emulating a function of an emulated processor is implemented, for example, in a "C" subroutine or driver, or some other method of providing a driver for the specific hardware as will be within the skill of those in the art after understanding the description of the preferred embodiment.

In FIG. 21, an example of an emulated host computer system **5092** is provided that emulates a host computer system **5000'** of a host architecture. In the emulated host computer system **5092**, the host processor (CPU) **5091** is an emulated host processor (or virtual host processor) and comprises an emulation processor **5093** having a different native instruction set architecture than that of the processor **5091** of the host computer **5000'**. The emulated host computer system **5092** has memory **5094** accessible to the emulation processor **5093**. In the example embodiment, the memory **5094** is partitioned into a host computer memory **5096** portion and an emulation routines **5097** portion. The host computer memory **5096** is available to programs of the emulated host computer **5092** according to host computer architecture. The emulation processor **5093** executes native instructions of an architected instruction set of an architecture other than that of the emulated processor **5091**, the native instructions obtained from emulation routines memory **5097**, and may access a host instruction for execution from a program in host computer memory **5096** by employing one or more instruction(s) obtained in a sequence & access/decode routine which may decode the host instruction(s) accessed to determine a native instruction execution routine for emulating the function of the host instruction accessed. Other facilities that are defined for the host computer system **5000'** architecture may be emulated by architected facilities routines, including such facilities as general

purpose registers, control registers, dynamic address translation and I/O subsystem support and processor cache, for example. The emulation routines may also take advantage of functions available in the emulation processor **5093** (such as general registers and dynamic translation of virtual addresses) to improve performance of the emulation routines. Special hardware and off-load engines may also be provided to assist the processor **5093** in emulating the function of the host computer **5000**.

In a further embodiment, one or more aspects relate to cloud computing. It is understood in advance that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

Characteristics are as follows:

On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.

Service Models are as follows:

Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models are as follows:

Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for loadbalancing between clouds).

A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure comprising a network of interconnected nodes.

Referring now to FIG. 22, a schematic of an example of a cloud computing node is shown. Cloud computing node **6010** is only one example of a suitable cloud computing node and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the invention described herein. Regardless, cloud computing node **6010** is capable of being implemented and/or performing any of the functionality set forth hereinabove.

In cloud computing node **6010** there is a computer system/server **6012**, which is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with computer system/server **6012** include, but are not limited to, personal computer systems, server computer systems, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud computing environments that include any of the above systems or devices, and the like.

Computer system/server **6012** may be described in the general context of computer system executable instructions, such as program modules, being executed by a computer

system. Generally, program modules may include routines, programs, objects, components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server **6012** may be practiced in distributed cloud computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed cloud computing environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

As shown in FIG. **22**, computer system/server **6012** in cloud computing node **6010** is shown in the form of a general-purpose computing device. The components of computer system/server **6012** may include, but are not limited to, one or more processors or processing units **6016**, a system memory **6028**, and a bus **6018** that couples various system components including system memory **6028** to processor **6016**.

Bus **6018** represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus.

Computer system/server **6012** typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server **6012**, and it includes both volatile and non-volatile media, removable and non-removable media.

System memory **6028** can include computer system readable media in the form of volatile memory, such as random access memory (RAM) **6030** and/or cache memory **6032**. Computer system/server **6012** may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system **6034** can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a “hard drive”). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a “floppy disk”), and an optical disk drive for reading from or writing to a removable, non-volatile optical disk such as a CD-ROM, DVD-ROM or other optical media can be provided. In such instances, each can be connected to bus **6018** by one or more data media interfaces. As will be further depicted and described below, memory **6028** may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of embodiments of the invention.

Program/utility **6040**, having a set (at least one) of program modules **6042**, may be stored in memory **6028** by way of example, and not limitation, as well as an operating system, one or more application programs, other program modules, and program data. Each of the operating system, one or more application programs, other program modules, and program data or some combination thereof, may include an implementation of a networking environment. Program modules **6042** generally carry out the functions and/or methodologies of embodiments of the invention as described herein.

Computer system/server **6012** may also communicate with one or more external devices **6014** such as a keyboard, a pointing device, a display **6024**, etc.; one or more devices

that enable a user to interact with computer system/server **6012**; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server **6012** to communicate with one or more other computing devices. Such communication can occur via Input/Output (I/O) interfaces **6022**. Still yet, computer system/server **6012** can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter **6020**. As depicted, network adapter **6020** communicates with the other components of computer system/server **6012** via bus **6018**. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server **6012**. Examples, include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

Referring now to FIG. **23**, illustrative cloud computing environment **6050** is depicted. As shown, cloud computing environment **6050** comprises one or more cloud computing nodes **6010** with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone **6054A**, desktop computer **6054B**, laptop computer **6054C**, and/or automobile computer system **6054N** may communicate. Nodes **6010** may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment **6050** to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices **6054A-N** shown in FIG. **23** are intended to be illustrative only and that computing nodes **6010** and cloud computing environment **6050** can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

Referring now to FIG. **24**, a set of functional abstraction layers provided by cloud computing environment **6050** (FIG. **21**) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. **24** are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

Hardware and software layer **6060** includes hardware and software components. Examples of hardware components include mainframes, in one example IBM® zSeries® systems; RISC (Reduced Instruction Set Computer) architecture based servers, in one example IBM pSeries® systems; IBM xSeries® systems; IBM BladeCenter® systems; storage devices; networks and networking components. Examples of software components include network application server software, in one example IBM Web Sphere® application server software; and database software, in one example IBM DB2® database software. (IBM, zSeries, pSeries, xSeries, BladeCenter, WebSphere, and DB2 are trademarks of International Business Machines Corporation registered in many jurisdictions worldwide).

Virtualization layer **6062** provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers; virtual storage; virtual networks, including virtual private networks; virtual applications and operating systems; and virtual clients.

In one example, management layer **6064** may provide the functions described below. Resource provisioning provides

dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and Pricing provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may comprise application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal provides access to the cloud computing environment for consumers and system administrators. Service level management provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

Workloads layer 6066 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation; software development and lifecycle management; virtual classroom education delivery; data analytics processing; and transaction processing.

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting. As used herein, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises” and/or “comprising”, when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below, if any, are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of one or more embodiments has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain various aspects and the practical application, and to enable others of ordinary skill in the art to understand various embodiments with various modifications as are suited to the particular use contemplated.

What is claimed is:

1. A computer-implemented method comprising:

initiating, by one or more processors, in a computing environment, an execution of a transaction by a processor, wherein the processor comprises a transactional core, and wherein the execution of the transaction is performed by the transactional core;

obtaining, by the one or more processors, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional transactional core in the computing environment, wherein the conflict is based on a request by the additional transactional core to the transactional core for data located in a cache of the transactional core during the execution;

determining, by the one or more processors, concurrent with the transaction, if the transactional core comprises an indicator, wherein the indicator is a super core indicator, wherein the super core indicator designates a super core status of the transactional core;

based on determining that the transactional core comprises the indicator, ignoring the conflict and utilizing, by the one or more processors, the transactional core to complete executing the transaction;

based on determining that the transactional core does not comprise the indicator:

aborting, by the one or more processors, the transaction performed by the transactional core;

incrementing, by the one or more processors, a value, wherein the value is associated with the transaction;

obtaining, by the one or more processors, the value;

evaluating, by the one or more processors, the value to determine if the value is equal to a threshold;

based on determining that the value is equal to the threshold, requesting, by the one or more processors, an indicator for the transactional core;

determining, by the one or more processors, that the super core status is available, wherein the determining if the super core status is available comprises:

accessing, by the one or more processors, transactional cores in the computing environment to determine if a threshold number of transactional cores comprise indicators; and

determining, by the one or more processors, that the super core status is available based on less than the threshold number of transactional cores comprising indicators; and

based on determining, modifying, by the one or more processors, the transactional core, to comprise the indicator.

2. The computer-implemented method of claim 1, wherein the data is located in the cache is marked as read data or write data.

3. The computer-implemented method of claim 1, wherein the ignoring further comprises:

denying, by the processor, the additional core access to the data.

4. The computer-implemented method of claim 1, further comprising:

queuing, by the processor, the request; and

based on completing execution of the transaction, fulfilling, by the processor, the request, by providing access to the data to the additional transactional core.

5. The computer-implemented method of claim 1, further comprising:

monitoring, by the processor, the executing of the transaction by the transactional core; and

based on the transactional core completing the executing by committing the transaction, clearing, by the processor, the indicator.

6. A computer-implemented method of claim 1, further comprising:

obtaining, by the processor, a request from the transactional core for the indicator.

7. The computer-implemented method of claim 1, wherein the threshold is a fixed value set by the program code or hard-coded in a software or a hardware element of the computing environment.

8. A computer program product comprising:
a computer readable storage medium readable by one or more processors in a computing environment, and

53

storing instructions for execution by the one or more processor for performing a method comprising:
 initiating, by the one or more processors, an execution of a transaction by a processor, wherein the processor comprises a transactional core, and wherein the execution of the transaction is performed by the transactional core;
 obtaining, by the one or more processors, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional transactional core in the computing environment, wherein the conflict is based on a request by the additional transactional core to the transactional core for data located in a cache of the transactional core during the execution;
 determining, by the one or more processors, concurrent with the transaction, if the transactional core comprises an indicator, wherein the indicator is a super core indicator, wherein the super core indicator designates a super core status of the transactional core;
 based on determining that the transactional core comprises the indicator, ignoring the conflict and utilizing, by the one or more processors, the transactional core to complete executing the transaction;
 based on determining that the transactional core does not comprise the indicator:
 aborting, by the one or more processors, the transaction performed by the transactional core;
 incrementing, by the one or more processors, a value, wherein the value is associated with the transaction;
 obtaining, by the one or more processors, the value;
 evaluating, by the one or more processors, the value to determine if the value is equal to a threshold;
 based on determining that the value is equal to the threshold, requesting, by the one or more processors, an indicator for the transactional core;
 determining, by the one or more processors, that the super core status is available, wherein the determining if the super core status is available comprises:
 accessing, by the one or more processors, transactional cores in the computing environment to determine if a threshold number of transactional cores comprise indicators; and
 determining, by the one or more processors, that the super core status is available based on less than the threshold number of transactional cores comprising indicators; and
 based on determining, modifying, by the one or more processors, the transactional core, to comprise the indicator.

9. The computer program product of claim **8**, wherein the data is located in the cache is marked as read data or write data.

10. The computer program product of claim **8**, wherein the ignoring further comprises:
 denying, by the processor, the additional core access to the data.

11. The computer program product of claim **8**, the method further comprising:
 queuing, by the processor, the request; and
 based on completing execution of the transaction, fulfilling, by the processor, the request, by providing access to the data to the additional transactional core.

12. The computer program product of claim **8**, the method further comprising:

54

monitoring, by the processor, the executing of the transaction by the transactional core; and
 based on the transactional core completing the executing by committing the transaction, clearing, by the processor, the indicator.

13. The computer program product of claim **8**, the method further comprising:
 obtaining, by the processor, a request from the transactional core for the indicator.

14. The computer program product of claim **8**, wherein the threshold is a fixed value set by the program code or hard-coded in a software or a hardware element of the computing environment.

15. A system comprising:
 a memory;
 one or more processors in a computing environment in communication with the memory; and
 program instructions executable by the one or more processors via the memory to perform a method, the method comprising:
 initiating, by the one or more processors, an execution of a transaction by a processor, wherein the processor comprises a transactional core, and wherein the execution of the transaction is performed by the transactional core;
 obtaining, by the one or more processors, concurrent with the execution of the transaction by the transactional core, an indication of a conflict between the transaction and at least one other transaction being executed by an additional transactional core in the computing environment, wherein the conflict is based on a request by the additional transactional core to the transactional core for data located in a cache of the transactional core during the execution;
 determining, by the one or more processors, concurrent with the transaction, if the transactional core comprises an indicator, wherein the indicator is a super core indicator, wherein the super core indicator designates a super core status of the transactional core;
 based on determining that the transactional core comprises the indicator, ignoring the conflict and utilizing, by the one or more processors, the transactional core to complete executing the transaction;
 based on determining that the transactional core does not comprise the indicator:
 aborting, by the one or more processors, the transaction performed by the transactional core;
 incrementing, by the one or more processors, a value, wherein the value is associated with the transaction;
 obtaining, by the one or more processors, the value;
 evaluating, by the one or more processors, the value to determine if the value is equal to a threshold;
 based on determining that the value is equal to the threshold, requesting, by the one or more processors, an indicator for the transactional core;
 determining, by the one or more processors, that the super core status is available, wherein the determining if the super core status is available comprises:
 accessing, by the one or more processors, transactional cores in the computing environment to determine if a threshold number of transactional cores comprise indicators; and

determining, by the one or more processors, that the super core status is available based on less than the threshold number of transactional cores comprising indicators; and

based on determining, modifying, by the one or more 5 processors, the transactional core, to comprise the indicator.

16. The system of claim **15**, wherein the threshold is a fixed value set by the program code or hard-coded in a software or a hardware element of the computing environ- 10 ment.

* * * * *